

100 Gbit Interconnects and Above: The Need for Speed

An OIDA Forum Report

Prepared by
Dr. W. S. Ring
WSR Optical Device Solutions LLC



© 2007 OIDA

Optoelectronics Industry Development Association
1133 Connecticut Avenue, NW - Suite 600
Washington, DC 20036
Tel: 202-785-4426; Fax: 202-785-4428
URL: <http://www.oida.org>

Report Documentation Page			Form Approved OMB No. 0704-0188		
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2007		2. REPORT TYPE		3. DATES COVERED 00-00-2007 to 00-00-2007	
4. TITLE AND SUBTITLE 100 Gbit Interconnects and Above: The Need for Speed				5a. CONTRACT NUMBER W911NF-06-1-0004	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Optoelectronics Industry Development Association (OIDA),1220 Connecticut Avenue, NW, Washington,DC,20036				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) US Army RDECOM ACQ Ctr -W911NF, 4300 Miami Blvd, Durham, NC, 27703				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES See also ADM002197.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 81	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Table of Contents

Table of Figures.....	iii
Tables	iv
1 Introduction.....	1
2 Industry status and issues.....	3
3 Background.....	5
3.1 Server market overview	5
3.2 Silicon industry technology drivers	6
3.3 Silicon industry companies	8
3.4 Communications industry overview	9
3.4.1 Optical components companies	11
3.5 Government initiatives.....	12
3.5.1 In the United States.....	13
3.5.2 In Japan.....	14
3.6 Protocols for the data center.....	15
3.6.1 Network introduction.....	15
3.6.2 Ethernet.....	16
3.6.3 Fibre Channel	19
3.6.4 InfiniBand.....	20
3.7 Standards development	21
4 Electrical interconnects overview	23
4.1 Overview of the interconnect segmentation.....	23
4.1.1 Intra-chip interconnects	23
4.1.2 Inter-chip interconnects	25
4.1.3 Backplane interconnects	25
4.2 RF considerations for backplanes and connectors	27
4.3 Channel model	28
4.4 ITRS roadmap for interconnects	30
4.5 3-D stacking and wafer level interconnects	30
5 Optical interconnects for backplane and boards	33
5.1 Fiber optic connectors and cable assemblies.....	34
5.2 Optical transceivers.....	36
5.2.1 Single channel – duplex.....	37
5.2.2 Parallel channel.....	38
5.3 On-board waveguides.....	41
6 Economics.....	47
6.1 Transceivers	48
6.2 Dilemma with the economics for optical component implementations.....	52
7 Breakout sessions.....	55
7.1 Session 1 summary.....	55
7.2 Session 2 summary.....	56
7.3 Session 3 summary.....	57
8 Roadmap	59

9	Recommendations.....	63
10	Summary and conclusions	65
	Appendix A – Forum Agenda	67
	Appendix B – Forum Attendee List	69

Table of Figures

Figure 1: Worldwide server revenue and market share for 2005 (US\$M)	5
Figure 2: Production linewidths in silicon (extended to 2010).....	6
Figure 3: Wafer size fabrication facility transition trend	7
Figure 4: Wafer facility cost trend.....	8
Figure 5: U.S. adoption of wireless handsets from 1990 to 2004.....	10
Figure 6: Overview of Japanese government programs related to optical interconnects	14
Figure 7: Simplistic view of the current network (Core, Metro, Access, and Last Mile)	15
Figure 8: Basics of the bridge/switched base Ethernet network (1994)	17
Figure 9: IEEE 802.3 reference model up to 10 Gbit/s data rates	17
Figure 10: Several basic Fibre Channel topologies	19
Figure 11: Fibre Channel protocol stack and overview of the translation.....	20
Figure 12: Example of the metal inter-connects layers in an IC	24
Figure 13: CPU and memory module example configuration.....	25
Figure 14: Backplane switching architecture	26
Figure 15: Different transmission line geometries	27
Figure 16: Example of a channel model.....	29
Figure 17: Example of a real channel model.....	29
Figure 18: Integration of memory methodologies being pursued today.....	31
Figure 19: Polaris 80 core and 3-D SRAM	31
Figure 20: Trend in data center fiber purchasing	35
Figure 21: Distribution of premise cable length sold to Corning customers	36
Figure 22: Examples of the duplex transceivers.....	38
Figure 23: Parallel optical link architecture using VCSEL arrays	39
Figure 24: High performance computing parallel transceiver solution.....	40
Figure 25: Conceptual overview of the optical PCB connections	42
Figure 26: Optical fibers in fibers in flex for waveguide in backplanes	43
Figure 27: Waveguide fabrication using polymeric materials.....	44
Figure 28: Mirror fabricated using a polymeric material	44
Figure 29: MT connectors for optical backplane connections.....	45
Figure 30: MT adapter to provide an optical connection from a PCB card	46
Figure 31: Distance vs. the cost per bit of transmitted data	48
Figure 32: Price erosion for Ethernet and Fibre Channel	49
Figure 33: Basic components of a SFP transceiver	50
Figure 34: Integration of memory methodologies being pursued today.....	59
Figure 35: Predicted optical bit rate evolution forecast for 2010.....	60
Figure 36: Evolution of optical components required for optical interconnects	61

Tables

Table 1: Top silicon semiconductor companies by revenue.....	9
Table 2: Review of current active component vendors and their financial results.....	12
Table 3: Summary of several DARPA-funded programs related to optical interconnects.....	13
Table 4: Basic view of the current protocol implementation with distance	16
Table 5: Examples of the aggregate rate, wavelength and fiber options	21
Table 6: Copper technology roadmap published by iNEMI.....	26
Table 7: Examples of material that can be used at very high frequencies.....	28
Table 8: Predicted year of production and process node line width.....	30
Table 9: Application table from the ITU, specifying link length classification	33
Table 10: Application table specifying transmitter and receiver requirements	33
Table 11: Ferrule dimension and connector types.....	34
Table 12: Fiber types in production and their over filled launch bandwidth.....	35
Table 13: Connectors initially developed for the SFF transceiver	37
Table 14: Cost estimates for a TO-46 header	51
Table 15: Estimated fabrication cost for a photodiode.....	51
Table 16: Generic cost per square foot for a clean room facility	52

1 Introduction

Today's electronic society is seeing increased demand for data transfer, internet downloads, online applications, video sharing, and storage. Banks, companies, universities, and governments require large secure data centers connected to secure networks. This continuously drives the development for more powerful servers and computer systems. Today's data centers are not single computer systems but networked systems with distributed processing, computational power, and memory. The systems are connected by 'nodes' which connect the pieces of the 'computer.' The configuration depends on the applications, functionality, and user requirements. This demanding environment is evolving and pushing for the development of higher speed interconnects.

Copper interconnects are currently used extensively within the data center environment. Fiber optic technology is used as a connectivity solution when a) higher performance is required and b) the cost differential compared to a copper solution is affordable. Fiber optic technology offers several key advantages over copper solutions for data transmission. Within today's data center, there is ubiquitous use of multi-mode fiber and optical transceivers for rack-to-rack connections. The fiber ports provide dedicated links with reach up to 300 meters on multi-mode fiber and 2 km on single-mode fiber. The fiber connections allow networks to connect to the wide area networks (WAN) or local area networks (LAN). As improvements have occurred in fiber optic transceiver performance and cost has been reduced, there has been further proliferation of fiber optics within the data center.

Traditionally, there has been a debate on the distance-cost crossover from a copper to an optic solution. As we delve into the interconnect market and technology, there are several motivations for optical implementation. To understand the industrial and commercial aspect, we look into the applications, history, market drivers, and future issues that are being discussed today.

OIDA held a one-day forum to increase the understanding on this topic and to look at high speed interconnects in the data center. The objective was to solicit input on the current issues related to optical interconnect technology which included examining the current research objectives of industry and universities and understanding the commercial objectives and maturity of today's technology.

Several organizations within the industry highlighted the technology issues facing copper interconnects for inter-chip, intra-chip, and board-to-board communications at a November 2004 OIDA forum. A reasonable question, considering the economic factors in the industry today, is: how can optical interconnects achieve greater penetration of the data center and server markets? This report reviews aspects of this subject, provides a synopsis of the information from the meeting, and presents several conclusions.

2 Industry status and issues

The computer and communications industries are interconnected today due in large part to the Internet, ubiquitous use of data, and transport of that data between servers, data centers, and access points. Today's computers are complex systems based on several technologies. The connection between the memory and the central processing unit (CPU) is not necessarily on the same board or in the same packaged integrated circuit (IC) chip. The data center is a distributed computer system with multiple nodes that have to be connected. These nodes sit at different levels, whether they are on the chip, in the server, in the rack, or at an alternative location. How these connections are made is highly dependent on the distance and the signaling speed.

The optical communications industry developed from the need for people to communicate. The telecom industry has been the driver for implementation of optical technology, which offers several advantages over copper wireline communications. As the computer industry moved to higher clock frequencies, smaller process nodes, and multiple core architectures, fundamental signaling issues drove the need to look at optics as an alternative architecture. Outside of the IC chip and microprocessor, the signal transport between chips and distributed nodes in the server/computer system are being reviewed for more extensive implementation of optical technology.

Copper media, multi-mode fiber, or single-mode fiber is used in today's data center to transport information between racks. The current architecture of the servers and data centers is based on copper interconnect technology. As the speed requirement between the different nodes increases, new approaches to the problem need to be developed. Copper interconnect technology continues to evolve and improve as the industry moves forward. The crossover point between optics and copper solutions has been debated and price and performance are regarded as key metrics. Signal integrity and thermal performance are additional metrics that cannot be ignored.

The computer industry has several organizations that look at the future bottlenecks that can impact the industry. For the semiconductor chip makers, the International Technology Roadmap for Semiconductors (ITRS) has developed several roadmaps and highlighted areas of concern. Interconnects is a key area. ITRS has highlighted a "red brick wall" that is looming for integrated circuit chip makers. Optical technology offers a solution, but the implementation and architecture to overcome the red brick wall are not understood. Relative to the traditional semiconductor industry, the optical component industry is immature. How the two merge to solve the problem remains to be seen. Does the optical component base need to be absorbed by the larger electronic semiconductor base of companies? Are there alternative hybrid implementations that can evolve in conjunction? Are they two different segments that will continue to stay apart or will there be a slow migration to optical implementation based on alternative optical technology currently in use today? In reality, copper interconnects will continue to be preferred until the level of maturity and implementation for short reach optical interconnects can meet several of the requirements that semiconductor chip companies expect today.

In response to the need for optical solutions to the current problems, several technologies are being investigated. For short reach links, vertical cavity surface-emitting laser (VCSEL) technology is ubiquitous today. The parallel VCSEL array transmitter and receiver (transceiver) is used today for server interlinks. Several government and industry initiatives have been investigating VCSELs for inter-board connections. The alternative approach currently receiving attention is “silicon photonics,” i.e., the integration of silicon IC technology and III-V light emitters. Silicon photonics is not a new concept: it has been investigated for more than 15 years and is utilized in production today in communications network equipment. The production of a group IV light emitting structure has not yet been realized; hybrid technology is believed to be the best route forward.

As we move to the intra-chip connection, the issues become more complex. Circuit board manufacturing and the reliance on FR-4 and its properties are a hindrance to new technology implementation and economies of scale for higher performance materials. Optical technology for circuit boards is developing along the path that provides waveguides imbedded in the circuit board. The removal of high speed electrical signaling traces needs to be complemented by changes in the input/output (I/O) architecture of the IC signaling pins and circuit card connectors.

In trying to understand the direction of the interconnect market and technology paths, there are several dimensions to the problem that need to be considered. Some of these are commercial and market driven, others are historical and experiential. The industrial direction, cost/performance, economy and scale of the silicon industry vs. the communications market need to be considered.

Economic issues have impacted the traditional optical transport business and the current optical components industry. This has led to a downturn in research and investment in optical component technology in the U.S. Globally, investment in optics is continuing. In Europe the 7th Framework program has established optical technology as a focus area. Both in Europe and Japan, investment in optics is now addressing broader market segments and applications. As Japan continues to implement its fiber-to-the-home (FTTH) program and broadband requirements, data center transport and signaling are being addressed. The U.S. companies researching this area need support from the government and a U.S. industry consensus.

3 Background

The computer industry and the optical communications industry have different requirements. The market is segmented into low cost, high volume PC/laptop sales, low-end servers, mid-range servers, and high-end servers. Each segment has different pricing and economic requirements. To look at the interconnect market we need to review several of the drivers.

3.1 Server market overview

The server/computer market is a multibillion dollar entity. The data centers utilize servers, data storage, and communication networks. The total server market is dominated by a few well know players. These companies utilize proprietary and/or standard protocols, develop their own chipsets, or partner with major chip manufacturers. The inside of the chassis is not interchangeable. The total estimated market for servers is around \$56 billion. This estimate includes cabling, processors, etc. The major players involved are IBM, Sun, Hewlett-Packard, Dell, and Fujitsu. Figure 1 shows the estimated share reported by International Data Corporation's (IDC) independent server market tracker.




Figure 1: Worldwide server revenue and market share for 2005 (US\$M)

(Source: IDC Quarterly server tracker)

The software which runs on the servers is broken down into three principle operating systems: Microsoft Server, Linux, and UNIX. The market is currently deploying 64-bit processors systems with multi-core designs. The industry has moved to multi-core designs rather than continuing clock frequency increases due to thermal and power

management issues. According to this tracker, blade shipments accounted for around 4.6% of the total market revenue in 2005. The blade architecture was reported as growing in popularity. Fiber optic cabling is utilized in the ‘high end’ servers. The main ongoing concerns for this market include power consumption and thermal management. The current I/O ports for the servers run at varying signaling rates. These range from 33 MHz (PCI²) to more than 2 Gb/s (PCI Express), with multi-drop busing used at the lower speeds and point to point links at the higher speeds.

3.2 Silicon industry technology drivers

The silicon industry is around a \$200 billion market. The major microprocessor and memory manufacturers employ state-of-the-art fabrication facilities. These large fabrication facilities are fully automated and can manufacture devices on 300 mm diameter wafers. The fabrication facility requirements depend on the market segment being served. The dynamic random access memory (DRAM), synchronous dynamic random access memory (SDRAM), and microprocessor segments are driven by constant reduction of device feature size to enable better performance and reduce costs. Figure 2 highlights this driver in linewidth reduction and processing node.




Figure 2: Production linewidths in silicon (extended to 2010)

The push in the microprocessor industry for smaller and smaller features has driven it to develop more precisely-controlled fabrication processes and implement different materi-

als, i.e., low-k dielectrics and copper conductors. The push for smaller feature size and larger wafers has made a dramatic impact on the fabrication facility investment cost. To remain competitive, many companies transition to larger wafer size. Figure 3 highlights the transition trend for wafer size production.




Figure 3: Wafer size fabrication transition trend

(Source: IC Insights)

The constant process node trend and wafer size change increases the investment cost for a semiconductor fabrication facility. Figure 4, highlights the increased cost in facility investment as the feature size has decreased.




Figure 4: Wafer facility cost trend

(Source: IC Insights)

The estimated cost of building a state-of-the-art 300 mm wafer facility is now close to \$3 billion. The primary manufacturer of microprocessors is Intel. It is the largest pure play silicon fabrication company and leads in investment in semiconductor technology.

The silicon industry has developed standard processes and design tool sets. The wafer foundry model used by several companies today was developed from a U.S. government sponsored program, Metal Oxide Semiconductor Implementation Service (MOSIS). The standard design sets and process tools enabled the silicon foundry model. There are several electronic silicon foundries, such as TSMC, Chartered Semiconductor, and United Microelectronics. These companies offer standard processes and tool sets at different process nodes. Typically, they are one generation behind the latest node being implemented by the vertically integrated player Intel.

For companies with revenue streams of around \$200 million or less annually, a fabless model makes tremendous sense and is preferred. It reduces the investment cost for these companies and allows them to develop new design and chip sets on essentially a shared cost line.

3.3 Silicon industry companies

The major semiconductor manufacturers utilize in-house fabrication, foundries, or a combination of both. The major companies in this field are outlined in Table 1. The table also highlights the major foundry suppliers. This list is based on 2005 revenue streams. The major semiconductor manufacturer is Intel, the principal CPU manufacturer.

Semiconductor Companies					
Company	Location	Revenue 2005	OI	Outsource	Notes
Intel Corporation	United States	\$ 38,826.00	\$ 8,664.00	No	
Texas Instruments Incorporated	United States	\$ 13,392.00	\$ 2,324.00	No	
AMD	United States	\$ 5,847.60	\$ 165.50		
Freescale Semiconductor	United States	\$ 5,843.00	\$ 563.00	Yes	
Infineon Technologies AG	Europe	\$ 10,060.30	\$ (340.00)	No	2006 income
STMicroElectronics	Europe	\$ 8,882.00	\$ 266.00	Yes	
Philips Semiconductors	Europe	\$ 7,452.90		Yes	
NEC Electronics Corporation	Japan	\$ 6,583.10	\$ 149.10	No	
Matsushita	Japan	-			
		-			
Silicon Semiconductor Foundries					
Company	Location	Revenue 2005		Outsource	
TSMC	Taiwan	\$ 8,103.60	\$ 2,292.70		
United Microelectronics Corporation	Taiwan	\$ 3,058.40	\$ (477.70)		
Chartered Semiconductor	Singapore	\$ 1,032.70	\$ (159.60)		

Table 1: Top silicon semiconductor companies by revenue

The semiconductor industry is cyclic in nature. The industry is driven by consumer demand, new product development and technology introductions. The optical interconnect requirements today are being driven primarily by companies that are producing micro-processors and memory chip sets.

3.4 Communications industry overview

The communications industry has evolved over the last 30 years as a result of significant changes in politics and global requirements. The initial development of communications networks was monopolized by the telecom operators. In the United States, this was AT&T, in the UK it was British Telecom, and in Japan, NTT.

During the early 1980s, politics had a major impact on the telecommunications industry in the U.S. and abroad. In 1984, the U.S. government decided to stimulate competition in the telephone sector. AT&T was broken up and 22 local Bell companies were formed into seven independent regional bell operating companies (RBOC), often called the “Baby Bells.”

In the 1990s, several significant events led to over-investment in the industry which resulted in the telecommunications “bubble.” The investment drivers in the communication networks came from the Internet and wireless telephone adoption. The lighting up of the dark fiber and resultant introduction of dense wavelength division multiplexing (DWDM) were major changes that occurred with the wireline fiber optic network. During the 1990s, most of the major equipment vendors spun out their optical component groups. This changed the business and research investment model for the industry in the United States and Europe. In Japan, the vertical nature remained due to close links between the major component suppliers and the system houses.

With the advent of the Internet, communication networks changed. The traditional telecom networks were based on voice traffic, with billing on a per minute basis. As the Internet expanded and proliferated, it became dominated by data traffic. This has filled the major network pipes today. Looking forward, the telecom carriers are now entering the video market and offer voice, video, and internet services (triple-play services). This has led recently to resurgence in demand for new network components. It is expected that video services will dominate fiber optic network traffic in the future. As high definition television (HD-TV) proliferates and on-demand video services are offered to consumers, more bandwidth will be consumed. New video download and sharing services, such as You Tube, are also filling the pipes. These new services are expected to impact the data center. For video-on-demand service, one potential solution to alleviate metro traffic buildup is local video caching.

Today, most family members have a cell phone. Figure 5 shows the adoption rate of wireless phones in the U.S. Cell phone operators such as Verizon, AT&T, Sprint, and T-Mobile have each built large cell networks and are now beginning to offer wireless fidelity (Wi-Fi) broadband access to the consumer and business user. This new service provides a roaming internet protocol (IP) network, with constant connectivity.




Figure 5: U.S. adoption of wireless handsets from 1990 to 2004

(Source: FCC)

With the breakup of the large vertical telecom companies, a horizontal optical component business model now exists in the United States. The current data center requirements and data transport services have changed the network structure. Due to its lower inherent cost

structure, Ethernet has started to become the network protocol of choice even for the telecom carriers.

Since the “telecommunications bubble,” there has been a switch in revenue leadership among storage, LAN, and telecom. Storage laser sales outstrip telecom laser sales today. This is understandable due to the proliferation of optical storage media and optical video players. The optical storage market for compact disc (CD), digital versatile disc (DVD), high definition DVD (HD-DVD), and Blu-Ray™, has different drivers than the optical communications market. The optical storage market is served by several chip manufacturers that provide both communications and storage devices. The storage market has a different set of requirements to enable read, write, and erase functions on the optical disc. The next generation optical storage medium is the Blu-Ray or HD-DVD disc. The optical heads now utilize three laser components based on the requirement for backward compatibility. The shortest wavelength used is 405 nm, which utilizes InGaN material to produce the source.

Meanwhile, the focus for optical communications companies remains developing faster and/or high power devices. The wavelengths of interest are centered on the transmission windows of the optical fiber, i.e., 850 nm, 1300 nm and 1550 nm. These companies are also engaged in developing faster I/O optical devices for servers and data centers, but remain cautious of the cost requirements for server links.

3.4.1 Optical components companies

Several companies serve both the communications and storage markets. The companies that are principally manufacturing devices for communications networks and I/O ports sell into the following application segments:

- Ethernet
- Fibre Channel
- SONET/OTN
- InfiniBand
- proprietary

The current optical components vendors have seen tremendous price erosion in their respective markets. With lower selling prices and high manufacturing costs, most companies have suffered significant losses since the “telecom bubble” burst. Table 2 highlights the operating income of several of the major component vendors.

Company	2006 Revenues	Operating Income (OI)	2004 Revenues	Operating Income (OI)	Notes
	(Millions)	(Millions)	(Millions)	(Millions)	
AVANEX	\$ 162.90	\$ (54.70)	\$ 106.90	\$ (124.10)	
Finisar	\$ 364.30	\$ (24.90)	\$ 185.60	\$ (113.80)	
Agilent			\$ 7,181.00	\$ 349.00	Note 1
Avago Tech	Private Company				
EMCORE	\$ 143.50	\$ 58.70	\$ 93.10	\$ (13.50)	Note 2
JDSU	\$ 1,204.30	\$ (151.20)	\$ 635.90	\$ (115.50)	
Cyoptics	Private Company		Private Company		
Bookham	\$ 231.60	\$ (87.50)	\$ 79.80	\$ (67.40)	
LuminentOIC			\$ 271.70	\$ (10.70)	
OCP	\$ 70.10	\$ 1.40	\$ 57.10	\$ (1.30)	
Eudyna	-	-	-	-	
Excelight					
Sigma-Links					
Opnext	\$ 151.70	\$ (30.50)	\$ 79.40	\$ (80.50)	
Mitsubishi	\$ 41,048.30		\$ 31,917.80	\$ 424.40	
Note 1: Agilent sold the Optics piece and it was formed as Avago Technologies					
Note 2: Include Sale of GelCore and Electronic materials					

Table 2: Review of current active component vendors and their financial results

The movement to transceiver modules compared to the traditional “golden box” discrete component has commoditized the market. There are several factors that have improved the operating income of several of the fiber optic transceiver companies, including the resurgence in the volume of sales to near-2000 levels. The optical components companies have also been actively engaged in downsizing and either pushing manufacturing off-shore or to outsource production. Significant changes in the company structures have provided an ability to improve gross margins and get closer to breakeven. Each company should be reviewed based on its core competency and product portfolio. Both Finisar and Opnext have shown evidence of profitability in several recent quarters. Generally, the sector is improving but is still not healthy enough to allow increased spending on R&D, a critical issue for the next generation of networks and component development.

3.5 Government initiatives

Governments around the world enable development of new technology and research through nationally funded programs. These programs impact the competitiveness and economic prosperity of a country. The U.S. government has traditionally funded research in areas that impact peoples’ lives through several of its agencies. This section highlights some of the facts related to research and development funding and information for the super computer segment and its impact on interconnect development in the United States and Japan. The principal reason to highlight these two countries’ programs is the continued race to achieve the highest performance super computer. This race spins off of development activities that impact the mid-range and low-end computer markets which benefit the overall computer industry.

3.5.1 In the United States

The U.S. government has been the primary driver for high performance computer development over the last 40 years. The impetus for this has been:

1. National security seeds
 - a. Intelligence
 - b. Conventional and nuclear arms
 - c. Weather forecasting
2. U.S. competitiveness vs. the rest of the world
 - a. Science and industry
3. Use of high performance computing (HPC) technology
 - a. Medicine
 - b. Climate research
 - c. Biology
 - d. Chemistry
 - e. Materials
 - f. Basic computer simulation (CS) tools research

This focus on HPC has enabled new advanced optical interconnect devices. There are several government agencies that fund research, including Department of Energy (DOE), National Science Foundation (NSF), National Security Agency (NSA), National Aeronautics and Space Administration (NASA), and the Defense Advanced Research Projects Agency (DARPA).

DARPA has provided a major focus for the computer companies by funding programs between industry, government, and universities. The level of interest and investment by the U.S. government in optical interconnect technology can be understood from several programs in which DARPA has invested. A number of these programs are highlighted in Table 3.

DARPA Project Name	Period	Budget (\$MM)
Analog Optical Signal Processing	2002-2005	37
Chip to Chip Optical Interconnects	2003-2007	45
Chip Scale WDM	2002-2005	40
Data in Optical Domain- Network	2002-2006	60
Optical CDMA	2003-2007	45
Photonic A/D Technology	1998-2001	40
Note: Reference Science and Technology Trends No.20 July 2006		

Table 3: Summary of several DARPA-funded programs related to optical interconnects

(Source: Science and Technology Review, Quarterly Review No. 20 - July 2006)

The continued focus on next generation and enabling technology by DARPA and other government agencies has helped the U.S. economy and has enabled technology leadership for many U.S.-based companies.

3.5.2 In Japan

The Japanese government has continued to invest heavily in optical communication components. The size of the fiber-to-the-home (FTTH) market is a clear indication of the demand within Japan and the government's goal for a connected society. When we look at computer systems and servers there has always been a race between the U.S. and Japan for the "biggest and best" super computer. The race remains heated, with the U.S. and Japanese governments investing heavily in this area.

The Japanese optical components companies are researching several key areas in long wavelength communications for high data rate applications. For example, companies are currently investing in directly modulated 40 Gbit distributed feedback (DFB) lasers. Several key results were presented at the Optical Fiber Conference (OFC) in 2007. There are a number of agencies in Japan that invest in new technology. A summary of the programs funded by the Japanese government (Figure 6) underscores the importance of optical technology to the country.

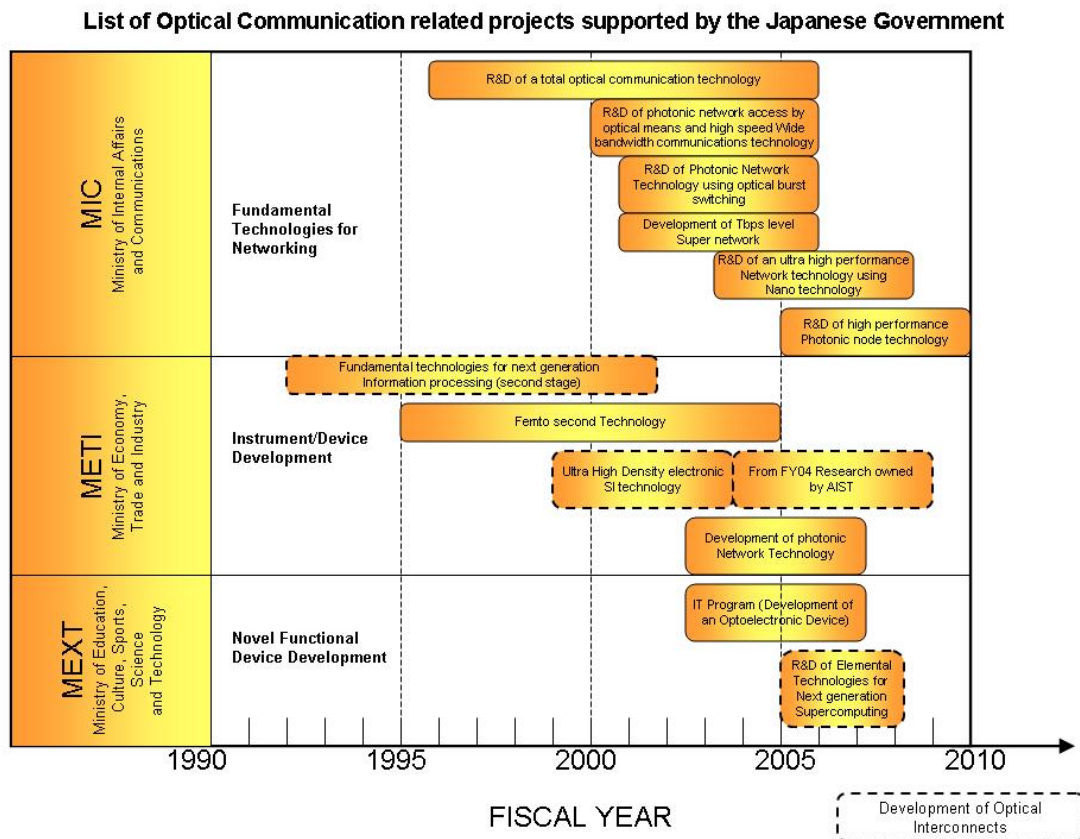


Figure 6: Overview of Japanese government programs related to optical interconnects
 (Source: Science and Technology Quarterly Review No. 20 - July 2006)

Within Japan's Advanced Industrial Science and Technology (AIST) group, development of optical backplanes is a key focus area. Some results from this group are presented later in this report.

3.6 *Protocols for the data center*

This section reviews some of the different protocols within the network and discusses the development of the different protocols and their current direction.

3.6.1 Network introduction

The current network infrastructure is still segmented by distance, application, and protocol. The standards bodies define the technology implementations and specifications. Data centers, server farms, and high performance computer centers exist to manage today's information society. These centers connect multiple servers, CPUs, memory, and storage devices. Each segment of the network relies on the three key technologies of optical, electrical, or radio frequency (RF). Which one is implemented is dependent on the transport requirement. A simple view of the communications network structure is given in Figure 7.




Figure 7: Simplistic view of the current network (Core, Metro, Access, and Last Mile)

Transport over the core, metro, and access is by internet protocol (IP). As we move into the data centers, the network structure becomes further segmented. The topology of the connected devices uses the same basic network configurations of rings, switches, and hubs, but the requirements for transactions, computation, and storage are different. Within the data center, several different transport protocols co-exist. These include Ethernet, Fibre Channel, InfiniBand, IP and, proprietary protocols.

These protocols essentially compete with one another for management of the data center network, but have different latency and CPU overhead requirements. The data center is an amalgam of several types of sub-system. The data center requires fast access memory, central processor units, storage (hard drives/tape drives), etc. It is essentially a 'big' computer providing computational efficiency, data access, and network functions. The

data centers today manage tremendous amounts of information. They must be reliable, provide redundancy, and have minimum down time.

There are several types of functions the data center can perform, dependent on the configuration. In a high-performance computer system, the center allocates resources and manages itself to appear as a single computer environment. The alternative is the implementation of multiple applications across several platforms with dynamic allocation.

The protocols implemented in the data center network are dependent on the link length, latency, and physical requirements. The following table provides an overview of some of the implementation choices in use today:

	Server to Server		Card to Card	Intra-Card
Length	10 - 300m	1-10m	0.3 - 1m	0.1 - 0.3m
No. of lines per link	1	1-10	1-10	1-100's
No. lines per system	10's	10-1000's	10-1000's	10-1000's
Standards	LAN/SAN Ethernet Fiber Channel Infiniband	Design Specific LAN/SAN Ethernet Fiber Channel Infiniband	Design Specific Ethernet - Infiniband PCI	Design Specific
Fiber Optic Transmission	Yes	Yes	?	?

Table 4: Basic view of the current protocol implementation with distance

There are several different technologies available to store data in the data center. These include magnetic tape, hard drives, optical media, and standard computer memory. The data center manager has the choice of using direct attached storage (DAS), network attached storage (NAS), or a storage area network (SAN). The type of interconnect used for data storage is discussed in a later section.

3.6.2 Ethernet

Ethernet is becoming the most pervasive network technology for communications. It is used in the data center, local area networks, and wide area networks. There is currently a strong movement for Ethernet to become the protocol of choice for telecom service providers as they upgrade their network infrastructure.

The Ethernet protocol is defined in layer 2 of the open system interconnect (OSI) model. Its functions are defined and developed by the Institute of Electrical and Electronic Engineers (IEEE). The original Ethernet LAN architecture has evolved since its inception in 1973. The original Ethernet network used a half duplex technology with carrier sensed multiple access/collision detection (CSMA/CD). As it developed, it moved to duplex operation and enabled the removal of collision detection, which can cause network slow-down. During the 1990s, the move to a switched-based architecture expanded the availability of Ethernet. A basic example is shown in Figure 8.

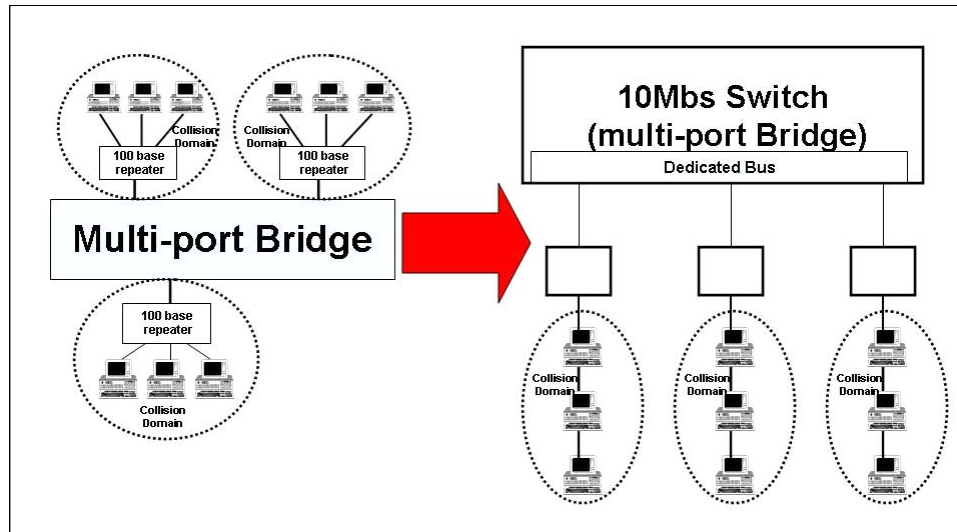


Figure 8: Basics of the bridge/switched base Ethernet network (1994)

The level 2 of the OSI model is expanded within Ethernet at the data link and physical layers. The architecture in the data link layer at level 2 is subdivided in the logical link control (LLC) and the media access control (MAC). This is illustrated in Figure 9.

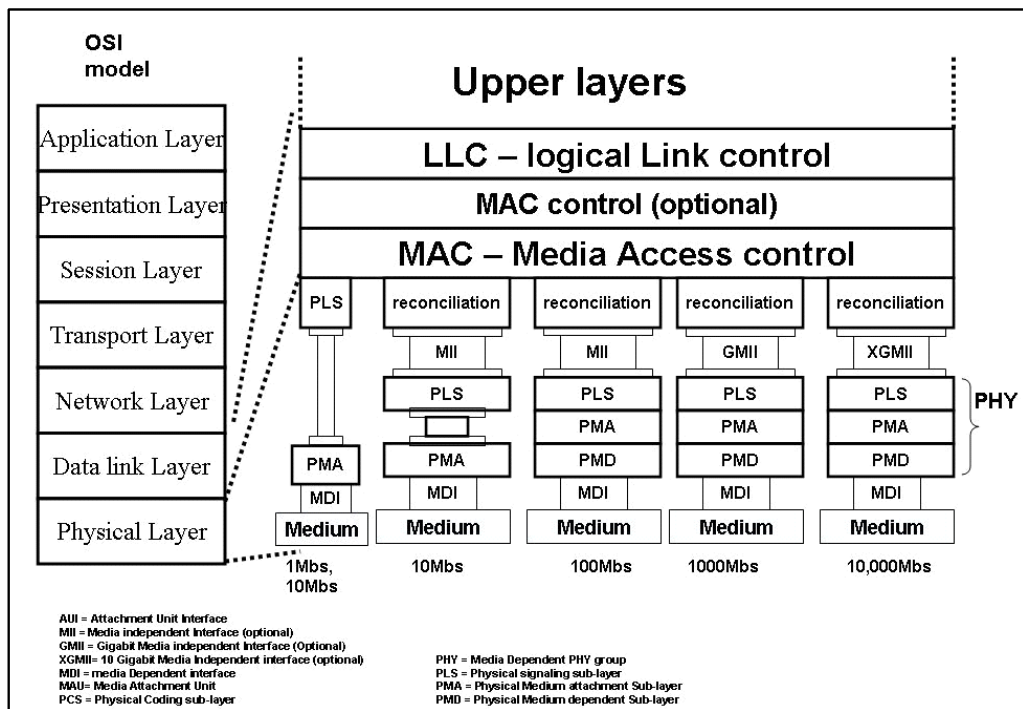


Figure 9: IEEE 802.3 reference model up to 10 Gbit/s data rates

(Source IEEE 802.3 Standards)

The LLC defines the common interface to the network layer, i.e., it provides interface points called service access points (SAP). The MAC layer and the physical layer implement the media specific functions. For example, the MAC has direct communication with a computer network interface card.

Today, Ethernet has expanded to the service provider backbone by using “carrier grade” Ethernet. The Metro Ethernet Forum (MEF), initiated in June 2001, has enabled the movement of Ethernet to the provider network. The Forum has worked on several topics to enable deployment. Changes were required as the carriers like determinism while the enterprise customer likes plug and play. To enable this, several topics on scalability have been, and continue to be, addressed. A carrier operator is typically looking for the following items:

- Open and standard interface
- Simplified network architecture
- Scalability
- Reliability
- Enhanced charging functions
- Enhanced security
- Quality of service (QoS)

The early Ethernet network was a best effort delivery system. The Ethernet has evolved to now provide ways to manage information differently for different user classes. The MEF defined specifications for Ethernet so that it can act like a service. The MEF certifies the equipment produced by its partners and members which enables deployment against a common standard. This helps the carriers deliver the service.

The next generation Ethernet network requires a provider backbone bridge as defined by 802.1ah. In this specification, each 802.1ah level encapsulates frames with a new MAC address and a new service tag. The nesting level summarizes the MAC addresses of the lower level with a backbone MAC address. The address summary allows a high degree of scaling without creating a MAC look-up table explosion. The end result is that Ethernet deployment in the Metro has become a realistic and reliable service solution.

Ethernet is also deployed extensively in the data center. The main issue with Ethernet in the server environment is the CPU overhead required to manage the connections. This is important for some applications and data center structures, but less important for others. To reduce the server CPU utilization, most Ethernet network interface cards (NIC) utilize transmission control protocol (TCP) offload engines. Essentially, the NIC card has a dedicated processor and off loads some of the communications overhead from the server itself. A concern in the high performance cluster environment is latency. Latency can be measured as the time from when a bit enters a switch to when it leaves the switch. Due to the nature of the packet size and routing functions, Ethernet has issues with pass through for high performance computing environments compared to the other protocols that are available.

Ethernet is in mass deployment in the enterprise environment at 1 Gbit/s and increasingly at 10 Gbit/s. The enterprise market is now looking at the next generation of routing requirements. The typical distance of interest for Ethernet solutions is less than 2 km and/or 10 km. Service providers are looking to Ethernet as the lowest cost solution for their networks. Ethernet has traditionally increased its data rate in increments of 10x the previous data rate. The 802.3 high speed study group is currently reviewing the next MAC rate within the IEEE standards development organizations.

3.6.3 Fibre Channel

The Fibre Channel (FC) protocol started development in 1988 with American National Standards Institute (ANSI) approval in 1994. The protocol and standard was developed to provide a seamless storage and backup capability. The Fibre channel protocol allows servers and workstations to connect directly to disk arrays, tape drives, or storage media. The protocol is a dedicated link approach to allow fast access and redundancy across a storage network. Typically, the FC network allows large data transfer with low overhead latency switching and minimal interruptions to data flow. It is different from Ethernet as it uses credit flow for point to point links. The basic idea is that the FC unit connects to another unit and logs on. The links then provide a maximum flow that can be transmitted on the dedicated link. Communication occurs until the flow is full. The FC network can be connected directly to the LAN or, today, even a WAN. This allows users a seamless gateway to their networked attached storage devices or servers. Figure 10 provides an overview of three Fibre Channel topologies.

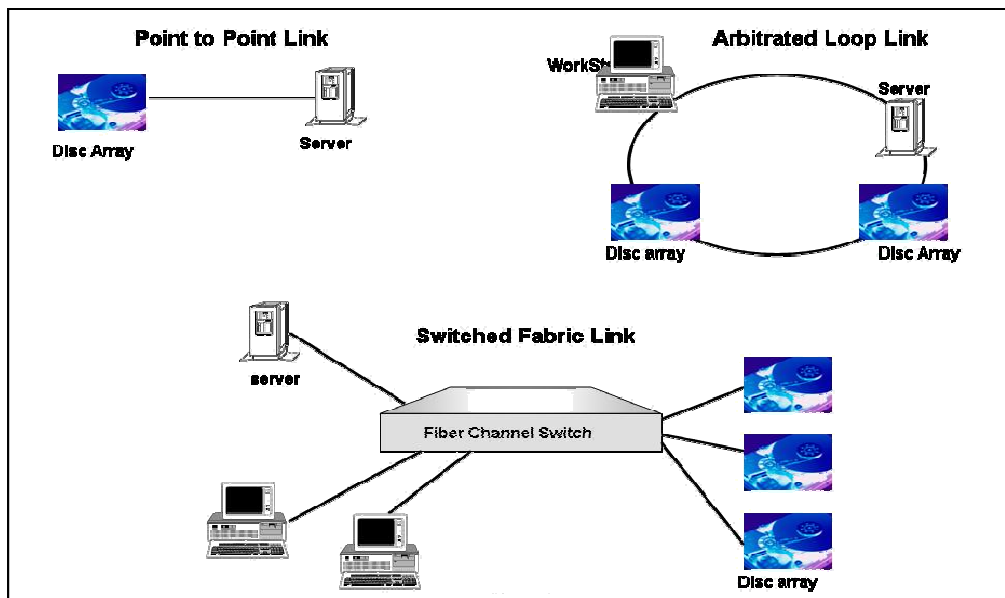


Figure 10: Several basic Fibre Channel topologies

Fibre Channel allows for redundancy and multiple storage capability. To connect to alternative devices or network, FC utilizes translation devices. The protocol allows

encapsulation of network and device protocols for transmission across the FC fabric. An overview of the translation devices is shown in Figure 11.

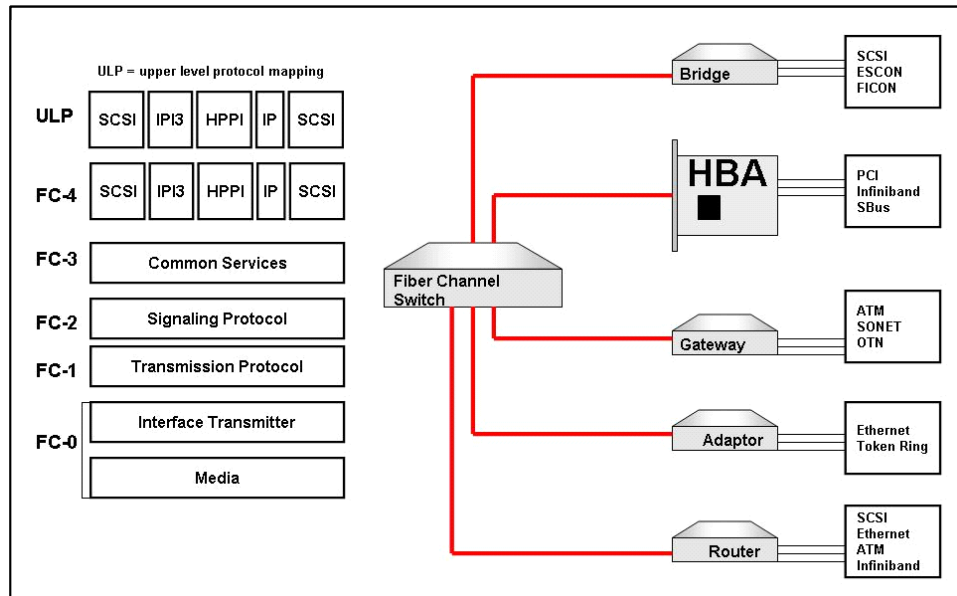


Figure 11: Fibre Channel protocol stack and overview of the translation

The host bus adaptor (HBA) is typically used as the node to connect to a server or storage device. These typically employ multi-mode fiber optic transceivers to communicate to the switch or port. The development of Fibre Channel is usually at 2x the previous data rate. During the optical bubble, a 10 G Fibre Channel standard was developed. This has seen minimal implementation. Instead the industry stepped back and developed the 4xFC standard. Currently, the T11 group is working on the 8.5 Gbit/s FC standard based on the 10 G FC links.

3.6.4 InfiniBand

With the movement in the data center to more of a cluster environment, low latency fast interconnects are required. The objective is simply to allow more efficient use of the resources in the data center. Cluster computing requires data to move at a rate between computers similar to the rate on the board. InfiniBand was initially developed to provide a server-to-server connection. It focused on employing fast connection speeds and delivering the lowest possible latency (ranging from 1.8 usec to 8 usec). The additional requirement for InfiniBand was to reduce the performance overhead placed on the server by managing the connection.

Several of the major companies in the server space are supporters of the InfiniBand standard. The principle benefits that it is advertised as providing are:

- Increased bandwidth fabric interconnects
- Low latency
- Enhanced density and reduced cabling requirements
- Multiple levels of redundancy

InfiniBand is still a switched fabric network and allows a point to point solution, similar to Ethernet and Fibre Channel. The architecture uses a single or multi-lane approach for the transmission and receiving of data. Some of the defined options for aggregate rate and wavelength are illustrated in Table 5.

InfiniBand Fiber Attachment Options				
Attachment Option	Data Rate (Gb/s)	Wavelength (nm)	Fiber type	Range (Meters)
IB-1X-SX	2.5	850	50/125um	2-250
			62.5/125um	2-125
IB-1X-LX	2.5	1310	SM	2-10,000
IB-4X-SX	10	850	50/125um	2-125
			62.5/125um	2-75
IB-12X-SX	30	850	50/125um	2-125
			62.5/125um	2-75

Table 5: Examples of the aggregate rate, wavelength and fiber options

It is believed that more fiber links will be required for server I/O interconnects as data rates increase. With the low latency achieved by Infiniband, there has been recent emphasis on the advantages of InfiniBand storage for blade server systems. Whether this will increase market share for InfiniBand will depend on the individual companies managing their data center.

3.7 Standards development

Standards play a very important role in the development of communication networks and components. There are several different standards bodies for the different levels of network architecture, including:

- OIF – Optical Internetworking Forum
- IEEE – Institute of Electrical and Electronic Engineers
- IETF – Internet Engineering Task Force
- ITU – International Telecommunications Union
- TIA – Telecom Industry Association

The use of standards, different protocol requirements, and technical specifications allows open network configuration and compatibility. Each organization has a vested interest in

its own requirements and development. The ITU is trying to understand the next generation telecommunications network. IEEE develops standard for local area networks and is currently addressing 100 Gbit Ethernet as the next generation data rate for switched networks.

Today's data center and servers utilize several of the standards developed by these bodies. How the implementation of protocols changes within the data center over the next decade is not clear. As we look forward to board-to-board and intra-chip optical interconnects, different standards bodies with less experience in optical technology may be involved.

4 Electrical interconnects overview

Within the data center, there are several different types of interconnects. To enable some distinction between the different requirements for optical technology, this report segregates the connections into intra-chip, inter-chip, and backplane interconnects.

Today, the majority of connections are copper based. Optical solutions can enable higher performance and fast data rates. The penalty that will be paid as signals move around optically is associated with the electrical to optical conversion either at the transmitter or the receiver. The extent of this penalty is connection-dependent. It needs to be analyzed as part of any cost or improvement driver.

4.1 Overview of the interconnect segmentation

The next few sections provide a basic overview of the types of interconnects and challenges associated with each.

4.1.1 Intra-chip interconnects

Within the IC, the individual transistor and circuit blocks must communicate via metal wires. As the photolithography line width decreases, the transistors become smaller and faster. This leads to more transistors per unit area and denser intra-chip interconnects.

The IC wiring hierarchy typically uses shorter wires nearer to the silicon surface and increasing longer wires at higher layers. The lower levels of interconnect are thin and are used in local routing. Intermediate layers are of medium thickness and used for semi-global routing. Finally, the top layers are the thickest and are used for global routing across the IC. The different levels of interconnect are customarily laid out in orthogonal directions to minimize cross-talk between adjacent levels. Furthermore, this convention helps to simplify routing patterns. Figure 12 provides an example of this wiring hierarchy.




Figure 12: Example of the metal inter-connects layers in an IC

As the clock frequency increases and the processing node decreases, global interconnects which span the chip exhibit higher resistance-capacitance (RC) time constants. This leads to increasing interconnect delay, transition time, and cross talk. In the majority of complementary metal-oxide semiconductor (CMOS) processes used today, the global interconnects are typically broken up into shorter links and signal is regenerated with repeaters. This breaks the long wires into smaller short segments which decreases the overall delay. The delay now becomes linear in nature rather than quadratic. The issue with the repeater is the delay introduced by the repeater itself and the associated power penalty because of this additional circuitry. For processors operating at Gigahertz clock frequencies, the global wires must be designed assuming resistance, inductance, and capacitance (RLC), not just resistance and capacitance.

The two main parameters characterizing the electrical interconnects are propagation delay and the interconnect bandwidth density. The delay can be reduced by increasing the interconnect width, but this is at the expense of a smaller bandwidth density. As the global interconnects are delay limited, it is expected that optical interconnects should be able to provide advantages here over the current approach. The current figure of merit for delay optimized global electrical interconnects is 1mW/mm, but this is expected to increase in the future. As the process node is lowered, the electrical bandwidth density can increase. This provides a moving target for optical interconnects to hit. For optical intra-chip interconnects, the main challenges are:

- technology implementation (silicon photonics, organic waveguides?)
- thermal and power consumption

- architecture and design rule changes
- package design

4.1.2 Inter-chip interconnects

There is a requirement on the circuit card to transfer data more efficiently between the CPU and off chip memory. Chip packages today connect to off chip memory using ball grid array packages that connect to the circuit card traces. The interconnection point and trace can provide loss of signal integrity. To implement an optical solution, we would need to change the I/O from optical to electrical, as shown in Figure 13.

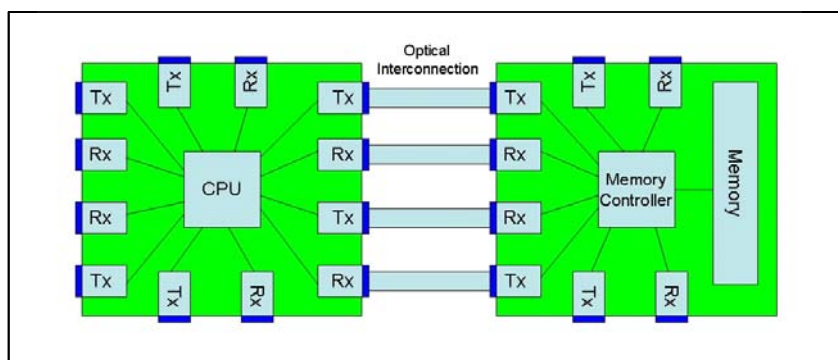


Figure 13: CPU and memory module example configuration

This represents a significant departure from current IC packaging and circuit design. There are fundamentally two elements that would need to change:

1. circuit card design and materials
2. I/O from the IC package

The design of this “OE extension” module that connects to the silicon circuit is not currently defined by any industry standard. There are several technical publications that show the potential of the VCSEL array to connect to a receiver array on board. The desire to provide the OE connection is a continuing challenge with the current industry dynamics.

The main implementation in practice today is rack-to-rack connections with parallel optical modules based on SNAP-12 or POP 4 multi-source agreements.

4.1.3 Backplane interconnects

Today’s servers are mainframes, racks, or blade server configurations. Copper backplanes connect between server blades or cards. The copper connectors in the backplane are produced by several key manufacturers, such as Tyco, Teradyne, Molex, and FCI.

These manufacturers have released a roadmap for copper connectors up to 40 Gbit/s (Table 6).

	1Gb/s	2.5Gb/s	3.125Gb/s	4.0Gb/s	6.125Gb/s	10Gb/s	40Gb/s
Connector	Teradyne connectors	HDM	VHDM & VHDM L-Series		VHDM-HSD	GBX	GBX
	ERNI Connectors	ERNet	ER met		Ermet ZD	ERmetZD	Ermet Zero XT
	Tyco Connectors	HM-ZD	Z-Pack HM-ZD		Z-Pack HM-ZD		
			MultiGig RT-1			MultiGig RT-2	MultiGig RT-3
Technology	FCI	Metral 2000	AIRMAX VS Metral 4000		AIRMAX VS		AIRMAX
	Fujitsu		FCN-261200x			FCN260D	
	Winchester		xcell				SIP-1000 I platform
	Molex	Molex is teradyne licensee					
Published by iNEMI							

Table 6: Copper technology roadmap published by International Electronics Manufacturing Initiative (iNEMI)

(Source: www.inemi.org)

The backplane provides for board-to-board interconnects. The architecture is a switched fabric architecture. There are principally two switch fabric architectures, as highlighted in Figure 14.

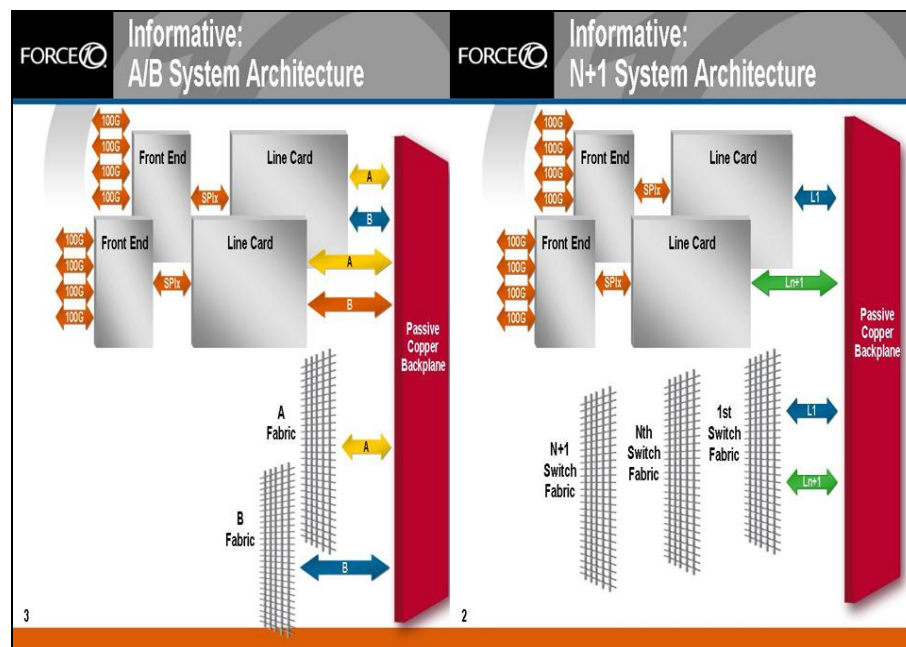


Figure 14: Backplane switching architecture

(Courtesy of J. Goergan, Force10 Networks – OIDA 100 Gb Ethernet Forum)

To move to an optical backplane, optical waveguides and connectors in the printed circuit board (PCB) need to be developed. Implementation would need to be standardized by the IPC-Association Connecting Electronics Industries or other standards organization. Several companies are developing optical waveguide technology for the PCB with great success, but several key challenges remain, including:

- architecture
- fiber/polymer or embedded waveguides
- implementation approach
- connector type

Optical will eventually be required. The market and application may not necessarily be in servers first, but could be in other markets such as cell phones.

4.2 RF considerations for backplanes and connectors

The delivery of signals between processors, memory, or backplanes today remain the principle domain of the electrical interconnect. The main issues for an increased data rate using electrical connectors are electromagnetic interference (EMI) performance, noise, attenuation, signal integrity, and cross-talk. Because of these issues, the prospect of using fiber optic solutions is more appealing. For electrical signals, the impedance of the path is dependent on the dielectric material constants, mechanical construction, and conductor materials. The simplest way to think of the connector is as a transmission line. Reflections need to be minimized to ensure that the signal remains “clean.” This requires careful control of all aspects of the connector mechanics and design. For example, a plated through hole (PTH) is typically viewed as a capacitive lumped element. Above 2.5 Gbit/s, the impact of the through hole must be viewed as a transmission line with a shunt element (or stub). This complicates matters as the plated through hole must be analyzed as a microwave structure.

To improve the signal integrity of copper cables and connectors, microwave design rules and modeling must be implemented. The simplest approach is to use microwave strip lines and transmission line theory. Reflections and losses must be minimized to provide good signal integrity. There are multiple approaches for transmission lines on the circuit board. Several useful transmission line examples are shown in Figure 15.




Figure 15: Different transmission line geometries

The interaction of the transmission line is dependent on several factors. One critical factor is the dielectric loss. Table 7 highlights different circuit card substrate materials that can be used at high frequencies.

Material/Process	Microwave Printed Circuit	Thin Film	Thick Film	Co-fired Glass Ceramic (LTCC)	Cofired Ceramic (HTCC)
Base Substrate	PTFE Glass Fiber PTFE Ceramic Polyester Glass Hydro carbon ceramic	Al ₂ O ₃ AlN BeO Quartz Glass/Ceramic	Al ₂ O ₃ AlN BeO	NA	NA
Conductors	Cu	Au, Al, Cu	Au, PtAu Ag, PdAg Cu	Au Ag PdAgCu	W Mo
Dielectrics	N/A	SiO ₂ Polyimide BCB	Glass Ceramics	Glass Ceramic	Ceramic Al ₂ O ₃ tape
Resistors	N/A	NiCr TaN	RuO ₂ Doped Glass	RuO ₂ Doped Glass	NA

Table 7: Examples of material that can be used at very high frequencies

For mass production FR-4 material is primarily used in today's manufacturing process. Other materials like Rogers 4003 provide an order of magnitude improvement in dielectric constant control for RF applications. The cost of these ceramic materials has decreased through widespread adoption in the wireless cell phone market. To ensure that the signal integrity is maintained, different microwave models can simulate the signal traces and interactions. These include analog-to-digital converter (ADC), high-frequency structural simulator (HFSS), and lumped element models. The scattering parameters (S-parameters) for the connector are the simplest to measure and model. This allows one to understand the reflection and transmission characteristics of the connector/cable assembly. As we move to very high data rates, the connector and cable assembly need to be viewed as one system.

In the backplane, there are additional techniques that can be employed to improve the received signal. These include pre-emphasis, multi-level adaptive signaling, or clock and data recovery (CDR). CDR uses an electronic IC to recover the signal and retransmit it to the next connection point. With pre-emphasis, some of the frequencies of the signal are boosted so that the loss impairments transform the signal to produce a clean signal at the other end of the connector. These techniques are extending the ability of copper cables and connector to achieve higher data rates and transmission distances.

4.3 Channel model

For the backplane, a channel model helps provide specification requirements for manufacturers to meet. This type of informative model has been adopted by several standards bodies. It allows parameterization and design criteria for the channel to be defined. Figure 16 shows the basic overview of the channel model for a chassis connection.

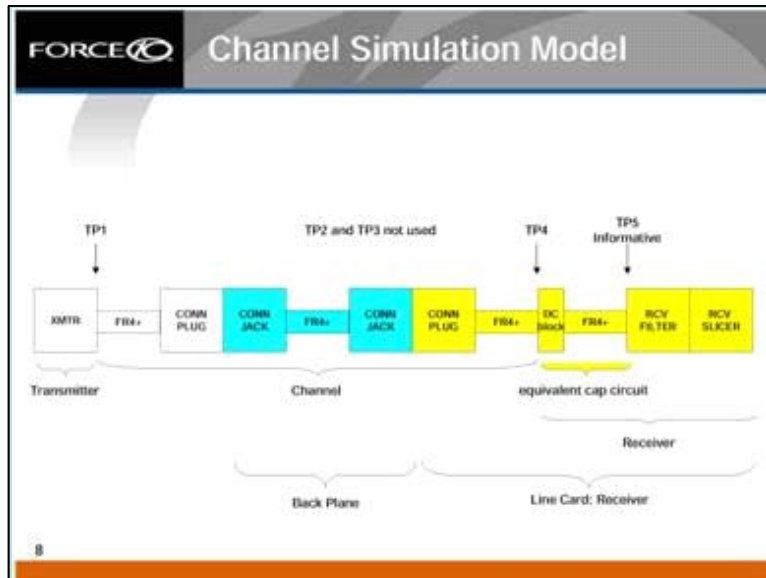


Figure 16: Example of a channel model

(Courtesy of J. Goergan, Force10 Networks – OIDA 100 Gb Ethernet Forum)

The trace length can be fixed or the maximum distance defined for the channel based on current circuit board material. An informative methodology can be employed which uses a fitting equation for signal loss. Figure 17 shows an example of the SDD21 of a channel model.




Figure 17: Example of a real channel model

(Courtesy of J. Goergan, Force10 Networks)

4.4 ITRS roadmap for interconnects

The International Technology Roadmap for Semiconductors (ITRS) has defined a path for the next generation of technology nodes (Table 8). As part of the road mapping process, the issues facing the semiconductor industry are discussed in this report.

Year of Production	2006	2007	2008	2009	2010	2011	2012	2013	2020
DRAM Stagger contacted metal 1 (M1) 1/2 Pitch (nm)	70	65	57	50	45	40	36	32		14
MPU/ASIC stagger contacted metal 1 (M1) 1/2 Pitch (nm)	78	68	59	52	45	40	36	32		14
Flash Uncontacted Poly Si 1/2 Pitch (nm)	64	57	51	45	40	36	32	28		13
MPU Printed gate length (nm)	48	42	38	34	30	27	24	21		9
MPU Physical Gate Length (nm)	28	25	23	20	18	16	14	13		6
Source: ITRS 2006 up-date										

Table 8: Predicted year of production and process node line width

(Source: ITRS Roadmap)

Several bottlenecks have been predicted. A number of these issues were reported at a November 2004 OIDA forum on interconnects. Two areas were highlighted that optical device technology might be able to address:

1. I/O pin count versus signaling rate per pin
2. intra-chip interconnects on the intermediate and global level (“red brick wall”)

These two areas are being addressed by research laboratories to find a potential solution. It is believed that continued advances in copper interconnect technology will dominate for the next 5 to 10 years until these bottlenecks can no longer be addressed by electrical solutions.

4.5 3-D stacking and wafer level interconnects

One current approach to providing a faster connection to memory on the chip is to stack die on top of one another. This approach reduces the interconnect length and allows higher speed interconnects. Currently, microprocessors are moving down the path of multi-core architectures. The advantage of the multi-core approach is the management of thermal load and heat sinking. As the implementation of on-chip optical interconnects remains a considerable challenge, the approach today is to extend the useful life of copper interconnects. It is currently believed that a resilient mesh approach will allow traditional electrical signaling to be used. As the processors which require memory and video cards become more process-intensive, the off-chip memory access will require high interconnection speed. There are two ideas that can be pursued to address this: die co-packaging or die stacking. Figure 18 highlights these approaches.

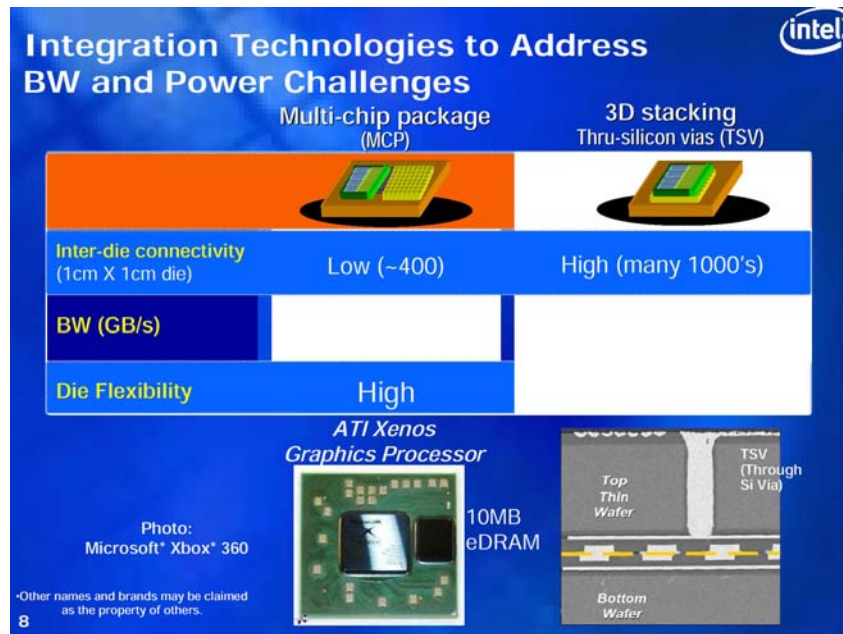


Figure 18: Integration of memory methodologies being pursued today
(Courtesy of J. Bautista, Intel Corporation)

The stacking of die using via technology for connection is being developed/implemented by several companies. In this approach, a via in a silicon spacer allows direct access to the memory chip from the CPU. Figure 19 shows an implementation for an 80 core processor under development at Intel Corporation.




Figure 19: Polaris 80 core and 3-D SRAM
(Courtesy of J. Bautista, Intel Corporation)

In this example, 80 SRAM tiles connect to the underlying IC circuit. This approach of 3-D stacking of die delays the need to provide an optical solution for off chip memory connection.

5 Optical interconnects for backplane and boards

There are several requirements for the optical interconnect to move further toward the chip domain in computer and server systems. The current technology in use today is based on the switch and router requirements for long distance optical networks. The dichotomy in the area is between short reach links (links < 300 meters in length) and long reach links (> 300 meters). The breakdown in optical networks is based on the link budget requirements and so has been segregated into transmitter and receiver type vs. distance and bit rate. A classic example is the SONET architecture. Tables 9 and 10 show the classification imposed on links for transport networks.

Application		Intra-Office			Inter- Office		
			Short Haul		Long Haul		
Source Nominal Wavelength nm		1310	1310	1550	1310	1550	
Type of Fiber		Rec. G652	Rec. G652	Rec. G652	Rec. G652	Rec. G652	Rec. G653
Distance (Km)		<2	15		40	80	
STM Level	STM-1	I-1	S-1.1	S-1.2	L-1.1	L-1.2	L-1.3
	STM-4	I-4	S-4.1	S-4.2	L-4.1	L-4.2	L-4.3
	STM-16	I-16	S-16.1	S-16.2	L-16.1	L-16.2	L-16.3

Table 9: Application table from the ITU, specifying link length classification

(Source: ITU-T G957 Specification)

		Unit	Values				
Digital Signal			STM-16 according to Recommendations G707 and G958				
Nominal Bit rate		kbit/s	2488320				
Application code			I-16	S-16.1	S-16.2	L-16.1	L-16.2
Operating Wavelength range		nm	1266-1360	1260-1360	1430-1580	1280-1335	1500-1580
Transmitter at reference point S							
Source type			MLM	SLM	SLM	SLM	SLM
Spectral Characteristics							
- Maximum RMS		nm	4				
-Maximum -20dB width		nm		1	<1	1	<1
-minimum side mode		dB		30	30	30	30
-suppression ratio							
Mean Launched Power							
- maximum		dBm	-3	0	0	3	3
-minimum		dBm	-10	-5	-5	-2	-2
Minimum Extinction ratio		dB	8.2	8.2	8.2	8.2	8.2

Table 10: Application table specifying transmitter and receiver requirements

(Source ITU-T G957 Specification)

With these standards specified for networks, optical components have moved to a module-based component structure. The modules are referred to as transceivers and contain memory, driver, and receiver chip sets, and can contain MUX and DE-MUX functions. The actual complexity of the module depends on the form factor and requirements from the customer of the switch, router, or host bus adaptor. The next section provides a review of the current types of transceivers. For server and board interconnects, the question on applicability may require alternative standard and reliability requirements. For intra-chip and inter-chip connections, the actual “module,” might develop into different

form factors than those readily available today. The main issue for the current entrenched supply base is the lack of incentives and investment money to exploit a potential high volume server chip interconnect market.

5.1 *Fiber optic connectors and cable assemblies*

The standard fiber optic components in servers and switches today are transceivers and fiber optic cable assemblies. These are based on well known and mass manufactured glass fiber. For circuit board routing, optical waveguide is an alternative device technology that can be employed. For Fibre Channel SANs and Ethernet LANs, fiber optic components utilize both single-mode and multi-mode fiber cable assemblies. These are typically duplex (two) fiber connections. For servers using proprietary and/or InfiniBand connections, parallel optic modules are utilized. Parallel modules use ribbon fiber cable assemblies. The different connectors for the fiber cable used today are MT-RJ, LC, SC and MPO connectors. All of these connectors provide a different ferrule size and fiber spacing. Table 11 provides dimensions of each connector type.

Connector Type	Ferrule Diameter	Fiber spacing	SM/MM
ST	2.5mm		Both
SC	2.5mm		Both
LC	1.25mm	6.25mm	Both
MT-RJ	2.5mm x 4.5mm	0.75mm	Both
MPO	3mm x 7mm	0.25mm	MM

Table 11: Ferrule dimension and connector types

For high density applications, the multi-fiber push on (MPO) fiber cable assembly is preferred. The cable assembly can be fanned out to other connector types or can simply be MPO-to-MPO. Alternative fiber management techniques have been developed where effectively the fiber is encased in a plastic jacket and routed with terminated connectors at either end of the assembly. This approach is simple but custom to the chassis being used. With the fiber connector technology, if the transceiver solution goes parallel, there are two choices:

1. Combine multiple wavelengths down a single fiber – e.g., D-WDM/CWDM
2. Utilize parallel arrays, e.g., 2-D and/or 1-D array ribbon fiber

Fiber optic cable is used extensively to transmit data across communication networks. It was first developed for the telecom network in the early 1970s. Improvements in drawing of the fiber and reduced impurities decreased the attenuation of single-mode fiber significantly. As fiber optic networks have evolved and new devices developed, the advantage of fiber has been employed on an ever increasing scale. Local area networks employed multi-mode fiber with LED-based transmitters for 100 Mbit/ links and today 40 Gbit links are being employed in the core networks of the telecom carriers. Fiber optic cables for I/O connections for servers have been growing in popularity as the benefits of the optic components increased and the reliability improved. Multi-mode fiber is used extensively for

short reach links because of the benefits of signal launch tolerance, cleaning, and low cost of the fiber optic transceivers.

There are different types of multi-mode fiber. The differences are characterized by the over-filled launch bandwidth. Due to its construction, the bandwidth of the fiber depends on the core diameter and material effective index profile. Table 12 summarizes the multi-mode fiber currently available.

Fiber	Wavelength	Core Diameter	Over Filled Launch Bandwidth
OM-1	850nm	62.5um	200
	1310nm		500
OM-2	850nm	50um	500
	1310nm		500
OM-3	850nm	50um	2000
	1310nm		
OM-3+	850nm	50um	4700
	1310nm		

Table 12: Fiber types in production and their over filled launch bandwidth

The multi-mode fiber modal properties have been studied extensively in the IEEE standards groups to allow higher data rate transmission over 62.5 um and 50 um core fiber. At 10 Gbit/s, electrical compensation techniques for the optical signal are used when the bandwidth of the fiber is too low. The single emitter standard developed for this application is the IEEE 802.3aq 10 GBASE-LRM standard. For server and SAN applications, these restrictions are not typical as new cable can be deployed. Multi-mode fiber is currently the fiber of choice for the data center.




Figure 20: Trend in data center fiber purchasing
(Courtesy of R. Grzybowski, Corning)

The high bandwidth fiber allows longer distances or higher data rate for equivalent distance. From the trend in cable type, higher bandwidth fiber is preferred for new installations over smaller bandwidth OM-2 or OM-1 fiber. Fiber length is also important to understand in terms of fiber cable interconnects. From the sales of cable length, Figure 21 shows clearly that 60% of the cable sold today is for distances below 100 meters.




Figure 21: Distribution of premise cable length sold to Corning customers
(Courtesy of R. Grzybowski, Corning)

This implies that 100 Gbit parallel optic transceiver links can meet a 100 meter distance requirement by utilizing OM-3 fiber.

5.2 Optical transceivers

Fiber optic transceivers are utilized in most network switches and SAN environments. For switches, the preferred transceiver is currently the pluggable module. For host bus adaptor cards, the through-hole version is preferred. The transceivers typically utilize InP devices for long wavelength 1310 nm applications and GaAs VCSEL devices for short wavelength 850 nm multi-mode applications. The parallel optical modules utilize VCSEL arrays for the SNAP-12 and POP-4 designs. The transmitter and receiver devices are the core of the transceiver. The InP and GaAs material systems today have been shown to allow modulation from the p-n junction device up to data rates of 40 Gbit/s.

The long wavelength devices are dominated by ridge and buried heterostructure devices. These are typically distributed-feedback (DFB) or Fabry-Perot (FP) laser devices that are either directly modulated or externally modulated. The cost of these devices is volume and overhead dependent. These devices are principally used in single-mode fiber applications. Single-mode fiber imposes a tight alignment tolerance on the packaging unless mode expanders or waveguide manipulation are utilized. They are predominately used for long reach links greater than 300 meters.

The two principal transceivers module interfaces are duplex (LC/SC) or MPO connector. Alternative interfaces for transceivers are used in the automobile market (plastic optical fiber), but these are not discussed in this report.

5.2.1 Single channel – duplex

There are multiple types of fiber connectors. The transceivers that are shipped in volume today are mainly LC duplex interfaces. When the market moved from the 1x9 footprint to the small form factor footprint, there were several interfaces/connectors initially developed (Table 13).

Connector	Ferrule size	Fiber spacing	Latch
LC	1.25mm	6.25mm	RJ-top 2 latch
MT-RJ	2.5 x 4.4mm	0.75mm	RJ-top latch
SC-DC	2.5mm	0.75mm	SC push-pull
VF-45	None	4.5mm	RJ-top latch

Table 13: Connectors initially developed for the SFF transceiver

The issues of repeatability, insertion loss, back reflection, and mechanical stability were all investigated. The two connectors that prevailed were the LC and MT-RJ connector. The LC is most widely used today because of the compatibility with TO-can technology and package diameter. The LC connector mated well with TO-46 and TO-56 assembly techniques for single-mode and multi-mode optical subassemblies. The fiber spacing provided enough mechanical tolerance to package two TO cans side by side. Other ceramic package types that utilized LC interface exist and these were developed later to enable cooling, optical isolation, or high speed RF feed through.

The market is segregated into different segments. Both Ethernet and SANs are predominately served today by short wavelength transceivers. Meanwhile, D-WDM, C-WDM, and SONET/OTN are mainly served by long wavelength transceivers. The FTTH market is different. This market typically uses fiber pigtailed product on a single fiber output/input.

The short wavelength market is dominated by the 850 nm VCSEL device. It utilizes a Bragg reflector in the device structure to launch the light through the top surface of the device. The VCSEL is a multi-mode device and replaced the 850nm edge emitter technology in Fibre Channel applications. Short wavelength transceivers utilize multi-mode fiber interfaces/connectors. These VCSEL devices are widespread, dominating the 1 Gbit Ethernet optical transceiver market and SAN market. The ratio of short reach to long reach transceiver ports in the Ethernet market is 70:30 (850 nm:1310 nm) at 1 Gbit/s. This is expected to reverse at 10 GbE. For Ethernet, the SFP and X-type transceivers dominate. Examples of these transceivers are shown in Figure 22.




Figure 22: Examples of the duplex transceivers
 (Courtesy of W. Jiang, JDSU – OIDA 100 Gbit Ethernet Forum)

At the higher data rates of 10 Gbit/s, additional ICs were initially incorporated into the transceivers to provide signal re-timer or multiplexer/de-multiplexer functions. The current trend in the industry is to push these functions back into the SER-DES on the host board. This is expected to reduce the electrical component cost within the transceiver. The reduction in footprint allows more ports per blade to be incorporated, increasing the density of 10 Gbit duplex transceivers on a board.

Work is progressing in the current Ethernet environment on developing 100 Gbit transceivers. The exact footprint, power consumption, and interface are still not clear. Due to cost constraints and development budgets, the industry trend is to utilize current component technology. This is an issue the industry will need to address as demand for high speed components becomes more pressing over the next few years.

5.2.2 Parallel channel

The parallel fiber architecture for interconnects was developed initially to solve the computer interconnect problem. A majority of funding of the basic research to enable these transceivers was provided by DARPA. The basic approach is similar to the ribbon cable used in the printers and computers in the 1980s. Instead of parallel electrical connection, a parallel optical connection is used. The transceivers are E-O converters with parallel signal inputs (Figure 23). The transmission (optical) lines must be compensated for skew, but this method allows increasingly higher aggregate bandwidth to be provided, which is proportional to the fiber count.




Figure 23: Parallel optical link architecture using VCSEL arrays

(Courtesy of J. Ojha, Avago Technologies)

The parallel links take advantage of VCSEL array technology. The array (photodiode or VCSEL) sits on top of an alignment substrate or electrical IC. The actual design is dependent on the manufacturer. Only a few companies have been successful with VCSEL array transceiver products. When these modules were initially released during the optical bubble, modules could fail catastrophically. The reliability concerns for parallel transceivers stem from the infant mortality issues related to dark line defects (DLD).

In the parallel module certain designs require the VCSEL to sit on top of the electrical circuit. This design allows thermal transfer from the IC to the VCSEL array, increasing the junction temperature. As reliability and infant mortality are accelerated with increasing temperature and drive current, the removal of the heat and thermal dissipation requires careful design. The problems of junction temperature increases and heat sinking of parallel modules have been published by several companies.

As the GaAs material system suffers from DLD, all VCSEL devices must go through a burn-in to weed out the weak devices. Typically, the burn-in forces the DLD to move into the active region if it is present. To overcome this problem, Agilent Technologies (now Avago Technologies) released several publications on their 1x12 transceiver suggesting that moving to a longer wavelength alleviated this problem. By adding indium into the quantum well and shifting the wavelength to 980 nm, they showed that infant mortality could be reduced.

The advantage the parallel transceiver offers is that the aggregate transmission data rate can be 4 to 12 times higher than the current duplex transceivers. Additionally, a parallel transceiver can provide equivalent aggregate data rate to a duplex transceiver but utilize lower electrical data rates. This approach can alleviate some of the signal integrity challenges that are seen at 10 Gbit/s and above. Figure 24 provides an overview of a parallel module.




Figure 24: High performance computing parallel transceiver solution

(Courtesy of P. Pepeljugin, IBM – OIDA 100 Gbit Ethernet Forum)

The VCSEL bandwidth, ESD, and reliability are related to the active region diameter. Conference publications have demonstrated 20 Gbit/s modulation in the laboratory. If the VCSEL is unable to achieve higher bandwidth, do we need to look at alternative device technology for parallel modules? In the serial transmission space the EA modulator or MZ-modulator approaches work well up to 40 Gbit/s but perhaps for higher data rates a different device will be required.

Alternatives to the POP-4 and SNAP 12 transceiver modules have been developed. The Quad Small Form Factor module, which is based on the XFP mechanical package, has been released to the market. It utilizes an MPO fiber optic connector to allow parallel transmission from a pluggable port, similar to the current XFP duplex transceiver. The design takes advantage of the XFP surface mount connector, allowing migration to 4x10 Gbit or 12x10 Gbit/s module. With the current development of high speed surface mount connectors, these rates could potentially move to 4x25 Gbit/s or 12x25 Gbit/s per electrical lane. A 12x25 Gbit/s parallel module would provide a 300 Gbit/s interconnect solution.

For board-to-board optical interconnects, the transceiver packaging will need to change. Potential solutions include co-packaging within the IC package or implementation of a silicon photonics solution. One of today's solutions is to integrate 850 nm VCSEL technology into the ball grid array (BGA) package or as an 'add-on' to the current silicon IC package. Co-packaging an O-E converter within the IC package will enable on-board optical interconnectivity. Several groups are using VCSEL technology because of the low drive voltages and alignment tolerance to multi-mode fiber or optical waveguides. Array technology is relatively mature and demonstrations of 2-D arrays connecting to circuit boards have been produced. The question of burn-in of the array and packaging/alignment is similar to the parallel module mechanics. This approach is a packaging exercise and very high precision assembly oriented. To move to a lower cost implementation, wafer scale integration will need to occur.

The alternative solution is currently silicon photonics. This approach is being investigated by several companies, including Intel. The main question with a wafer scale approach is generating the light transmitter or light source. Several silicon photonic building blocks have been demonstrated including modulators, receivers, and transceivers. These demonstrations have used extrinsic sources. The simplest integration of the light source is currently achieved using standard flip chip die bonding techniques. To achieve a wafer scale source, the integration into the silicon fabrication line would be preferred. One method is to use wafer bonding of III-V material to the silicon. Wafer bonding is a common process in the LED industry. The idea is simply to add some III-V semiconductor that can generate the photons for the silicon laser resonator. This topic is covered in more detail in the OIDA report *Silicon Photonics: Challenges and Future*.

Integrated O-E converters would enable the optical PCB industry. The problem will be whether it should be based on 850 nm VCSELs or 1310 nm silicon photonics. The actual wavelength will need to be determined to enable the right properties of the waveguides to be implemented. 1310 nm VCSEL technology has been developed but is not in mass deployment. The technology remains relatively immature compared to today's 850 nm devices.

Fundamentally, the conversion to an optical signal should be part of the silicon IC to achieve the lowest cost. The next section addresses the state of play of optical PCB technology.

5.3 On-board waveguides

The basics premise of the optical printed circuit board is to allow transmission through or on the PCB with an optical guide that exhibits low loss and dispersion. There are several materials that can be utilized to provide the optical waveguide, including

- glass fiber
- polymers
- sol gels
- glass

The structure and dimensions of the waveguide are determined by the geometry, materials, and signal wavelength. As highlighted before, there are two distinct approaches being developed within the industry:

- VCSEL and parallel based transmitter technology
- Silicon photonic transmitters

The coupling of the transmitters to the optical circuit board is a key area that needs development. The approach needs to allow blind mating with standard pick and place equipment. Highly tolerant optical signal paths need to be developed between the waveguide and the IC's optical transmitter and receiver circuitry. Figure 25 provides a conceptual overview:




Figure 25: Conceptual overview of the optical PCB connections

(Courtesy of Peter Van Deale, University of Ghent)

The current backplanes in the server architecture utilize FR-4 circuit cards and electrical connectors. There has been a tremendous amount of effort in optical PCB card R&D. Optical technology reduces EMI emissions and works well with flexible topologies. Several university and industrial companies have been developing optical waveguides. A few of these approaches are highlighted below.

Method 1: Glass fibers in flex

The simplest approach is to use standard fiber. The fiber can be placed on a polymer or “plastic” backplane and then routed and spliced to enable the I/O required. The fiber is sealed within a polymer/plastic for rigidity and durability (Figure 26).




Figure 26: Optical fibers in polymer to form a flex cable for waveguide in backplanes

(Courtesy of Peter Van Deale, University of Ghent)

Method 2: Polymer waveguides

The concern with polymer waveguides, dating back to the first silicon bench developments, is stability. Polymers have been found to have unstable refractive index with exposure to moisture environments. Today, polymers have been developed that can provide a stable refractive index profile over time. These polymers can be fabricated into waveguides using standard photolithography (Figure 27) or other techniques, including:

- Gray-scale photolithography
- Molding
- Laser ablation

The losses for the some of these waveguides are < 0.1 dB/cm at 850 nm wavelength.

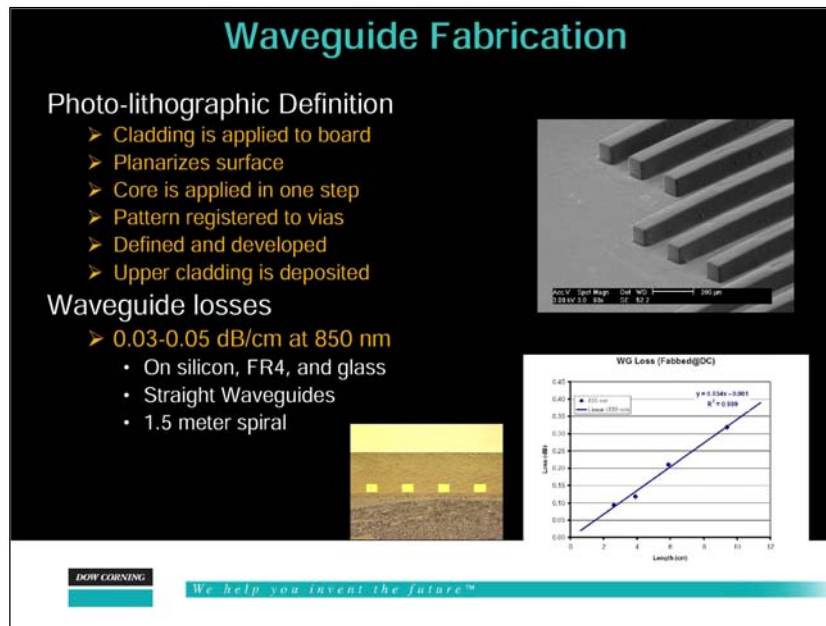


Figure 27: Waveguide fabrication using polymeric materials
 (Courtesy of J. V. DeGroot, Jr., Dow Corning)

Molding of the waveguides provides a road to mass production. Several researchers have published work where silicon masters are used to control the waveguide mold tolerance. Molding also offers the ability to form different structures. Mirrors can be incorporated into the waveguide layout and produce ultra smooth surfaces. Figure 28 provides an example of a mirror fabricated with a polymeric material.




Figure 28: Mirror fabricated using a polymeric material
 (Courtesy of J. V. DeGroot, Jr., Dow Corning)

PCB Waveguide connection

The MT connector is utilized in fiber optic cable assemblies, which makes it an ideal choice for the optical PCB and backplane. In Japan, a government program has been developing optical backplane technology to America's Carriers Telecommunications Association (ACTA) standards. The backplane connection utilizes polymers for encasing the fiber which are terminated with an MT connector. Figure 29 highlights the interconnection scheme under development.




Figure 29: MT connectors for optical backplane connections

(Courtesy of H. Itoh, AIST, Japan)

This approach enables high speed optical daughter board to backplane connection. The mechanical alignment and interface is well understood and available today. An alternative is to attach the MT connector to the polymer waveguide PCB (Figure 30). This is not as elegant as the glass fiber connection but it is a step forward for polymer-based waveguides.

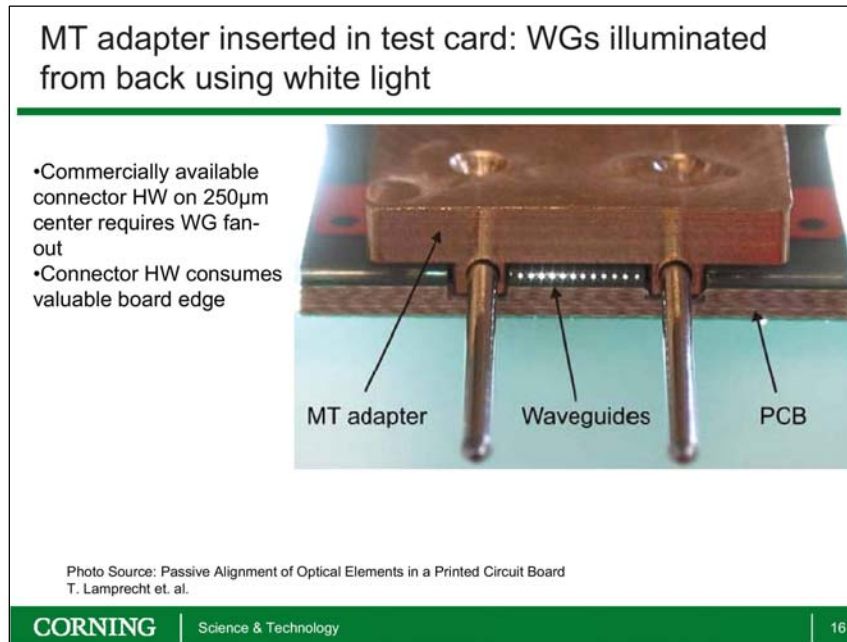


Figure 30: MT adapter to provide an optical connection from a PCB card
(Courtesy of R. Grzybowski: Corning)

Further research and development is required to push optical PCBs to the mainstream. The development of standards and optical IC front ends need to be pursued to enable mass deployment of this technology.

6 Economics

The economics for the implementation of improved I/O connections is dependent on the health of the supply chain and technology investment requirements. It is important to understand both the needs of the end user and the capabilities of the current suppliers.

From a historical look at the fiber optic and copper interconnects/cable market, it has been assumed that fiber will increase in deployment as the data rate increases. As new standards and data rates have been released, however, fiber is having difficulty displacing copper solutions. Each time the optics community feels that it will gain improved share, the electrical community provides a new or better solution. For example, at 10 Gbit Ethernet it was believed that a copper solution could not be achieved. Yet a 10 Gbit base-T standard has been developed and 100 Gbit Ethernet copper cable solutions have been proposed in the current IEEE HSSG study group. The copper connector and cable assembly companies are profitable while the fiber optic companies remain unprofitable or are beginning only now to break even. This disparity plays into the economic challenges for high speed I/O. Both the computing and data center markets have aggressive cost targets. This can directly impact the roadmap for fiber optic interconnects, based on the current economic issues prevailing in the component industry.

During the 100 Gbit Ethernet workshop held in August 2006 by OIDA, several companies noted that there is a real need to increase the I/O connection speed. This is becoming increasingly important in the high performance computer (HPC) market. There has been fiber optic deployment and it will migrate from the high end server to the low end servers as costs decrease. One of the key requirements to ensuring successful deployment is cost. The estimated cost requirement per bit of transmitted data for the HPC environment is distance driven (Figure 31).




Figure 31: Distance vs. the cost per bit of transmitted data

(Courtesy of P. Pepeljugoski, IBM – OIDA 100 Gbit Ethernet Forum)

The question of how to achieve this with optical module technology is challenging. The cost of the module, optical sub-components, and volume play a large role in the economics. The current cost today of a 1 Gbit short reach (SR) optical module is nearly equivalent to the price of Cat5e cable a consumer can buy from Staples...\$20 to \$30. For the 10 Gbit/s duplex modules, the price is an order of magnitude larger and currently in the \$300 and above range. The comparative cost depends on the application, i.e., short reach module or long reach transponder.

6.1 Transceivers

Duplex transceivers have had incredible price erosion over the last 10 years. As new modules have been introduced, the price premium for each new introduction has been eroded. The price expectation is dependent on the market segment. For example, the expected price of the 40 Gbit transponder is about 2 to 2.5 times that of the 10 G transponder price while the price for a 4.25 Gbit Fibre Channel transceiver is expected to be equivalent to the 2.125 Gbit Fibre Channel transceiver. The problem with this expectation is that the current transceiver vendors are struggling to be profitable. With no premium for new module introduction, there is very little money available to develop new transmitter and receiver technologies to enable the next generation of components. An exam-

ple of the price erosion based on the cost per Mb/s per linear inch of face plate was presented at the OIDA micropackaging workshop and is shown below:




Figure 32: Price erosion for Ethernet and Fibre Channel

(Courtesy of K. Jackson, EMCORE)

The Fibre Channel market is fortunate that most of the investment in 10 Gbit and 40 Gbit technology occurred during the 1990s, prior to the optical bubble bursting. As the enterprise market looks at a projected market for 100 Gbit Ethernet of only 100,000 units by 2010, the return on investment argument is no longer valid. As we move to these speeds, the fundamental technology that was developed through the 1990s has not seen any radical change. Hermetic packages are preferred and TO-cans remained a significant volume runner. The development of VCSEL technology has had a huge impact in the storage market, but long wavelength VCSEL technology has struggled to break ground. The promise of ubiquitous low cost 1310 nm VCSEL technology is currently competing with low cost FP and DFB laser technology. As the data rate requirements have increased, DFB and FP infrastructure, production, and yields have improved to enable lower manufacturing costs. As the data rates move to above 10 Gbit/s, the external (lithium niobate) or internal (electro-absorption) modulator technology is used. For the enterprise and telecom markets, the fiber non-linearity and impairments make the standard non-return-to-zero (NRZ) serial solution unattractive at 40 Gbit/s for long reach applications. Additionally, electronic driver technology is more expensive and large voltage swings are required. At lower data rates, low voltage swings and inexpensive IC technology provide a cost argument to implement parallel links over serial type solutions.

For current optical duplex transceivers, the construction and devices inside the ‘box’ provide significant cost challenges to meet a \$/Gbit metric. There are several factors that make this difficult. These include the module construction, type of semiconductor device, and manufacturing location. The basic components of a duplex transceiver are not complicated, as shown in Figure 33.




Figure 33: Basic components of a SFP transceiver

(Courtesy of B. Ring, WSR-ODS)

The cost of a transceiver can be broken down into several levels. At the macro level, the basic components are the electrical and optical subassemblies that are soldered together and put into the casing. The subcomponents, i.e., transmitter optical subassembly (TOSA) and receiver optical subassembly (ROSA), are a large percentage of the cost in 1 Gb/s modules. At higher data rates, the electronics are a larger part of the cost breakdown. This is one of the reasons that the SFP+ module is being developed.

The principle sales in the storage market are in VCSEL-based products. The main package used to produce the TOSA is the TO-can. The elements of the package are essentially the following:

- VCSEL die
- Monitor photodiode
- TO-can
- Lens cap
- ADM assembly

The package material price for the TO-can has been reduced significantly due to the volume market for CD lasers which has driven down manufacturing costs. In volume, the

package costs for these elements are much less than a U.S. dollar. An example of a TO-46 price in low volume is provided in Table 14.

Quantity	1000	5000	10000	50000
Description				
TO-46	\$ 1.00	\$ 0.45	\$ 0.40	\$ 0.29

Table 14: Cost estimates for a TO-46 header

(Private Communication)

The principle material cost driver can be the semiconductor device and its assembly processing cost. The cost of the device will depend on the company cost structure, including depreciation and operating costs. To provide information relevant to the debate, Table 15 provides a photodiode cost assumption based on a loaded labor rate of \$250/hour.

Die Cost model example						
Die Size		300μm x 300μm		Assumes 1 wafer per batch and approximated processing times. Loaded labor rate depends on fabrication facility and business model. Consumables cost ignored.		
Useable Wafer Area		40mm ²				
Expected Die per wafer (Unyielded)		17778				
Wafers/batch		1				
Loaded Labor Rate		\$ 250.00				
Grown EPI Wafer Cost		\$ 300.00				
Step	Operation	Yield	Cum Yield	Time (hrs)	Yielded Material Cost	Yielded Labor & OH
1	START WAFER					
2	DIELECTRIC DEPOSITION	100%	100%	2	\$ 300.00	\$500
9	PHOTO: DIFFUSION	100%	100%	6	\$ 300.00	\$1,500
23	ZINC DIFFUSION	90%	90%	5	\$ 333.33	\$1,389
30	DIELECTRIC DEPOSITION	100%	90%	2	\$ 333.33	\$500
33	PHOTO: P-CONTACT	100%	90%	3	\$ 333.33	\$750
47	PHOTO: P-METAL	100%	90%	3	\$ 333.33	\$750
58	"P" METAL DEPOSITION	100%	90%	3	\$ 333.33	\$750
66	WAFER THINNING	100%	90%	4	\$ 333.33	\$1,000
69	"N" METAL DEPOSITION	100%	90%	3	\$ 333.33	\$750
76	PROBE TEST	95%	86%	25	\$ 350.88	\$6,579
79	DIE SEPARATION	80%	68%	6	\$ 438.60	\$1,875
84	WAFER PICKING AND DIE INSPECTION	90%	62%	4	\$ 487.33	\$1,111
				66	\$ 4,210.14	\$17,454
			Total Cost/Wafer =	\$	21,664.08	
			Cost/die =	\$	1.22	

Table 15: Estimated fabrication cost for a photodiode

(Source: Cost Model: OIDA InP Foundry Workshop)

The example in Table 15 is loaded labor rate dependent. Several factors influence this including the processing linewidth. Assuming a linewidth of 1 µm requirement (i.e., active region width of a 1310 nm semiconductor laser), the fabrication investment would be several tens of millions of dollars for new equipment. The facility construction cost depends on the air handling requirements and clean room class. Table 16 shows the approximate cost per square foot for a fabrication facility vs. air handling requirement.

Criteria	Class Limits				
	10	100	1000	10000	100000
Air Changes per hr	600	300-480	150-250	60-120	Oct-40
HEPA Filter coverage %	100	70-100	30-60	10-30	5-10
CFM per Sq. Ft.	90	65-36	32-18	16-9	8-5
Typical Filter Efficiency	99.9997	99.997	99.997	99.997	99.97
Typical Filter Velocity	60-110 FPM	50-90 FPM	40-90 FPM	25-40 FPM	10-30 FPM
Air Flow type	Unidirectional	Unidirectional	Mixed	Mixed	Mixed
Typical Return Air Flow System	Raised Floor	Low Wall	Low Wall	Low Wall or Ceiling	Low Wall or Ceiling
Generic Cost Estimate (\$ per Sq Ft)	600-750	450-650	160-260	60-70	40-50

Table 16: Generic cost per square foot for a clean room facility

(Source: OIDA InP Foundry Report)

As the die size is reduced, the number of die per wafer increases. This should decrease die per wafer cost, but the loaded labor rates depend on the utilization of the facility, operating costs, and depreciation. For small market sizes requiring few wafers per year, the cost of ownership of the fabrication facility is questionable. For the VCSEL component, the additional wafer validation and burn-in costs are added to the die cost. Table 15 highlights that depending on the company overhead, the die can either be the most significant cost in the module or have a comparable cost to the outside material purchase cost of the TO-can.

With advances in semiconductor packaging and technology, alternative packaging and assembly is available to transceiver manufacturers. Principally, they could leverage the technology utilized in LED packaging or IC manufacturing to reduce costs. LED packaging today utilizes lead frames with high thermal dissipation. To move to lead frame packaging requires development of non-hermetic VCSELs and laser/photodiode devices. A non-hermetic environment can drastically increase the infant mortality rate of active optical semiconductor devices. This would necessitate moving to several of the techniques initially adopted in the silicon industry.

6.2 Dilemma with the economics for optical component implementations

For optical interconnects to become more ubiquitous in the computing arena, there are several challenges for the optical community, including:

- Cost
- Performance
- Reliability
- Size
- Thermal dissipation
- Interconnection interface

The primary targets for optical component implementation continues to be cost. Several companies from the computer and server markets provided their input. The targets provided to the community were:

- Board-to-board: \$1 Gbit
- Chip-to-chip: \$0.25 Gbit

The current parallel optic transceiver manufacturers stated that perhaps \$4 Gbit was achievable today. The prices are aggressive for current technology and packaging approaches. In addition, testing is seen as a major obstacle. Alternative schemes such as those being pursued using silicon photonics or integrated III-V OEIC may be able to meet these targets.

The server architecture today is not designed for optical interconnects. This would need to change for optics to make progress. The debate of the economics of a copper cable solution vs. a fiber optic solution is multifaceted. It will continue to progress, but slowly, as both economic factors change and technology advances. In the end, if there is a will there is a way.

7 Breakout sessions

There were three breakout discussions held at the workshop. The objective was to debate the issues the group had been presented during the day. The breakout sessions were led by three moderators who provided written feedback from their meetings to the group. The three questions debated were:

1. The system requirements for optical and electrical interconnects and the current bottlenecks: how will interconnects need to develop to address the increasing speed requirements?
2. How do the optical communications and silicon industries address the short reach interconnect issues?
3. What are the technology issues for next generation active/passive optical and copper interconnects? Where do we need to focus research?

7.1 Session 1 summary

Session 1: *The system requirements for optical and electrical interconnects and the current bottlenecks: how will interconnects need to develop to address the increasing speed requirements?*

1. There are a variety of systems from small to large, from few nodes to many nodes, servers, storage, memory, so you can't exactly draw a bright line where optics or electronics will or won't be used.
 - a. Can be extra value to optical connection, for instance orthogonal mating of boards that may drive a decision
2. Different applications
 - a. Telecom, not strained today
 - b. Server consolidation → drives bigger pipes
 - c. High performance computing → target communications cost 10-20% of system cost
 - i. System architecture driven by tradeoffs of memory, CPU power, number of nodes, comm cost (which goes as n^2 on nodes)
 - d. Data centers – conservative in technology adoption
3. Today: \$3 Gb would be attractive for optics; \$1 would be no-brainer Cu replacement
4. Power, latency, airflow/backplane, special value to optics for system reasons (e.g., security)
5. On board waveguides, very interesting and may be part of solution
 - a. No infrastructure, no one can build the boards for 5-7 years

- b. These kinds of developments may make possible new architectures that take advantage of new capabilities

Timeframe	2006	2008	2011	2016
Chip-chip (~10cm) ²	\$0.25/G 1Tb/s 7,8		\$0.02/G 10Tb/s ^{9,10,11}	\$0.002/G 100Tb/s ⁹
Board-board (~1m) ²	\$1/Gb 200Gb/ch ^{3,6}		\$0.5/G/b 1Tb	\$0.1/G 5Tb
Rack-rack (10-100m) ¹	\$5/Gb 20Gb/ch *Critical need today	\$3/G 100G	\$1/G 400G	\$0.25/G 1Tb/ch

1. This market driven primarily by price
2. These markets are driven by needs/performance advantages GIVEN that price is “reasonable”
3. Current bandwidths range up to 400 G today, nominal number
4. Price is the value at which optics becomes attractive but other points (power, etc.) will drive ultimate decision
5. Optics may add additional value due to unique capabilities so could be priced higher in some apps
6. Backplane cost dramatically increased for higher bandwidth by back drilling
7. Cell processor is .8 Tb/s today
8. Memory wall, application-specific integrated circuit (ASIC) pin count is driving need for new solutions
9. ITRS red brick wall...no solutions seem available
10. Fiber count as important as pin count, can't have hundreds/thousands of fibers/package, optics can enable WDM, highly muxed bw/fiber, is this the need for SiPhot?
11. At 10 Tb and \$0.02 G, optics cost is \$200...reasonable?

7.2 Session 2 summary

Session 2: ‘How do the optical communications and silicon industries address the short reach interconnect issues?’

The debate centered on the question of a real definition for short reach. For the working group, we defined this to mean inside the box or cabinet.

For inter-rack connections, the InfiniBand standard has growing acceptance from vendors. This is due to the low latency of the connection. The working group discussed many issues relevant to the question. Some of the key points raised included:

1. We need predictive failure for short reach interconnects
2. The system architecture today does not lend itself well to optical solutions
3. Silicon photonics is in its relative infancy compared to other III-V solutions
 - a. There is no real field demonstration of hybrid silicon photonics except for III-V light emitters bonded to silicon for the light source
 - b. The failure modes are not currently understood

There are key drivers in computers that are different to traditional optical communications links. The principle drivers are:

- volume
- cost
- \$ Gb
- \$ Watt

The current optical modules that are used for inter-rack communication utilize multi-mode parallel links. These are expensive but are forecast to achieve \$4 Gb. The target for optical modules was discussed as being \$1 Gb for inter-box connections and lower for inside the box. To reduce the cost of the transceiver, it was suggested that eliminating testing and providing relaxed specifications would enable cost reductions. One of the key concerns of the module manufacturers included connectors. Optical connectors are expensive, as is the cabling. Development of lower cost optical connections would enable more proliferation of optics in the data center.

Alternative solutions that could reduce this cost include moving away from the traditional transceiver to on-board optical interconnects. To drive this, the computer architects need to design for optics. This will enable a more optics-friendly environment.

One of the overriding guidelines that arose was that the optical module or on board process needed to be compatible with CMOS.

7.3 Session 3 summary

Session 3: What are the technology issues for next generation active/passive optical and copper interconnects? Where do we need to focus research?

This session focused on the technology issues for next generation requirements. Several points of debate were raised and the following points were discussed:

1. We need to provide more consideration for parallel architectures
2. Dealing with the skin effect for copper devices increases the longevity for interconnect applications
3. Is there a next generation connector for the optical fiber? Should the industry look to fiber connectivity at the fiber interface?

4. There needs to be a compatibility of O/E packaging with IC processing. (e.g., adhesive performance with reflow)

To understand the next generation requirement, we need to understand the next level of system architectures. The key issue is routing of the signals in the servers. Can we utilize wavelength division multiplexing (WDM), passive optics, or different network components? The group recognized that technology roadmaps for the process nodes and disciplines are essential.

When we look at the data requirements, the group asked itself, ‘What does it take to get to Terabit solutions?’ Is it WDM, single-mode fiber, or multi-mode fiber?

As the group moved forward, it asked for cost effective transmitter and receiver components. The price target for these was 1 Gb/\$. Most of the arguments centered on whether or not we can get to the low cost device if we have the volume. It was stated that this argument is similar to the CD/DVD market. The group thought that economies of scale could enable this by taking advantage of wafer technology.

The copper connector is familiar for computer and server manufacturers. The issues of signal integrity and skin effects are well understood. The copper companies are a small ‘group’ which understands the deficiencies and provides standardized product. To enable optics penetration into traditional copper markets in the computer industry requires process standardization, joint efforts between companies, and clear development roadmaps. It would be preferred if a single material component platform could be used.

Several of the optical component vendors suggested that packaging cost is a limiting factor. Fiber alignment, testing at each process step, increases the optical device/module cost. This needs to be understood to enable the cost requirements to be met. Leveraging a high volume application was recognized as the route to achieving this. In summary, the group recommended the following:

1. The loss needed to be reduced for electrical connectors. This would need different material and efforts to reduce cross-talk and EMI.
2. Cost needs to be reduced for the current optical technology for IO the packaging. This could be achieved by more integration of the optical components. The advantages of WDM and parallel transmission should be leveraged. The group requested that specific roadmaps for components be generated.
3. New technology needs to be implemented for the next generation of IO. A single material platform would enable this if it could leverage wafer scale methods.

8 Roadmap

One of the objectives of the Interconnect Forum was to draw conclusions and develop a path forward for review within the industry, i.e., a roadmap. Several organizations develop roadmaps in technology to provide potential paths forward. The computer industry has been following Moore's law for microprocessors and as the dimensions become smaller, new effects and roadblocks appear. The ITRS regularly reviews the current trend in IC semiconductor nodes and development of new technology. The silicon industry has adopted new approaches to resolve signaling across chips and to improve signal integrity. We have seen the implementation of copper and low-k dielectrics. As signaling requirements increase, 3-D stacking and silicon via technology to local memory chips has been developed. These methods continuously improve the performance of copper interconnects within chips and delay the implementation of O-E front ends at the output pins of the IC.

As we look forward, the microprocessor industry continues to advance and enable new applications. It is expected that terabit instructions per second will enable a new host of applications. The timeframe for this is the next 10 to 15 years. Figure 34 highlights the development objectives and potential applications moving forward.

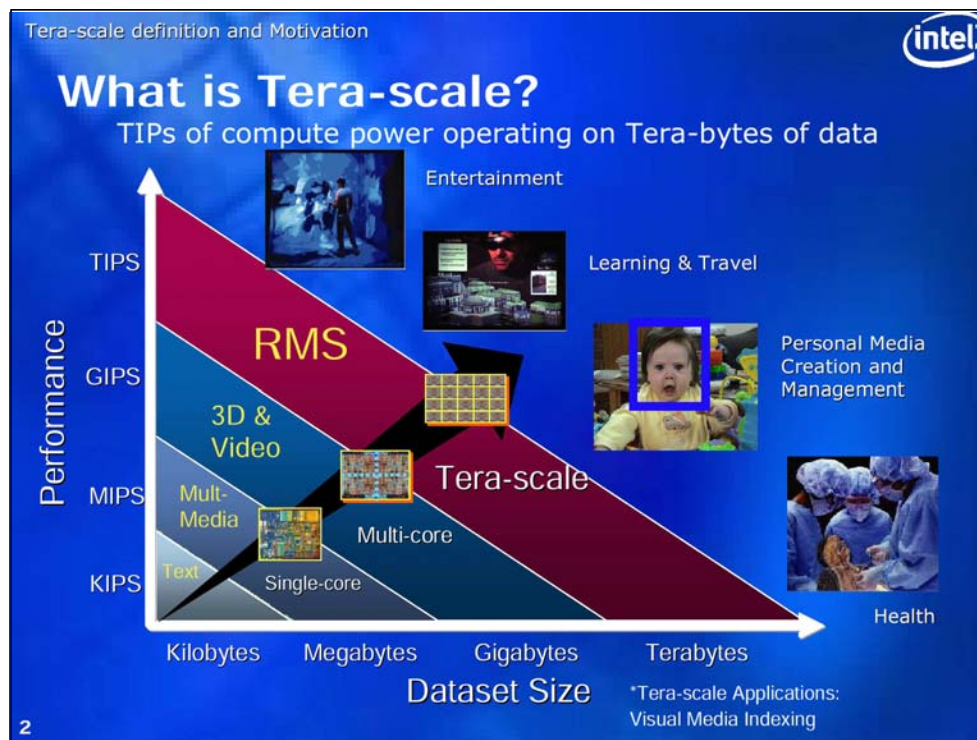


Figure 34: Integration of memory methodologies being pursued today

(Courtesy of J. Bautista, Intel Corporation)

As the industry pushes forward to multi-core processing and “tera-scale” architectures in the future, both the parallel interconnection using copper and fiber need continued research effort.

As we look at the traditional optical component development and communications roadmaps, it was expected that SONET/OTN would move to 160 Gbit/s as the next step after 40 Gbit/s. Figure 35 highlights the trend in optical communications data rate development.




Figure 35: Predicted optical bit rate evolution forecast for 2010

(Courtesy of J. Ojha, Avago Technologies)

Today we are seeing several developments within the industry in optical communications. The reversal of OTN vs. Ethernet leadership in data rate is occurring as carriers are aligning their transport protocols to carrier class Ethernet.

As we look forward to the next generation of devices, components, and technology, the economics of the industry directly impact the future development paths. Optical development of on-chip optical interconnects is expected to be delayed as silicon photonics is still immature and will require further development.

Looking to future expectations, the data rate is projected to increase and several bottlenecks currently occurring need to be resolved. The 2007 report from iNEMI suggests that optical component technology is too immature to provide optical interconnects even by 2017. For optical interconnects to impact board-to-board connections, the implementation

of the optical backplane will need to overcome alternative lower cost technology options. It is expected that optical backplanes will be required for HPC environments.

A few important “take-aways” from the meeting include:

- Optical technology needs to drive to lower cost and be compatible with the silicon industry process to enable introduction in a co-packaged fashion.
- The E-O conversion needs to be low power and low cost.
- Forum participants expect to see a push for high aggregate bandwidth either through a silicon photonic platform or a III-V semiconductor platform.

Figure 39 highlights several expectations drawn from the OIDA meeting.

Optical Interconnects




Figure 36: Evolution of optical components required for optical interconnects

9 Recommendations

Fiber optics offer clear advantages as we move to higher data rates within the computer industry. The question on how this is best resolved will require industry collaboration and greater interaction between the traditional optical community and the silicon semiconductor industry association. The silicon industry will need to address the interconnect roadblocks on the horizon by providing a change in the architecture from electrical to optical. Several approaches are being developed today to address the same problem. The following are some recommendations resulting from the many discussions at the workshop. The recommendations are a synopsis of several of the debates held during the breakout sessions.

- The server companies need to provide clear guidance on their fiber optic and optical circuit card technology requirements if they wish to enable adoption within the next 10 years.
- OIDA should engage with IPC and other optic standards organizations to discuss how to standardize optical PCBs and their optical interface.
- With both silicon photonics and VCSELs attempting to address the optical interconnect market, standardization on wavelength would help address this conundrum.
- The cost of an optical module for the SAN and data center will require a new packaging technology implementation. Leveraging the LED packaging technology should be considered by the fiber optic supply base.
- The copper connector industry has defined the roadmap for next generation connectors up to 40 Gbit/s. The fiber optic industry needs to understand how this fits into the transceiver module requirements and what additional connector technology should be developed in conjunction with module manufacturers.
- Current packaging of fiber optic transceivers has migrated to South East Asia; the U.S. government should recognize the lack of fiber optic packaging infrastructure in the U.S. and develop a program to improve American competitiveness in this area.

10 Summary and conclusions

The forum held by OIDA provided valuable insight in the current paths being developed for optical interconnects. Clearly, the use of copper interconnect for short links remains healthy and improvements in performance keep the implementation of optical connections from becoming mainstream.

Several different technology approaches were presented that are being developed. This was highlighted at the Silicon Photonics forum (held the next day). The discussions on architecture highlighted a crucial aspect that must be resolved within the computer industry for optics to advance. The computer architecture must evolve to favor optical over electrical interconnects. Currently, optics is utilized for high end rack-to-rack connections using parallel optical modules. These offer aggregate rates of 30 Gbit/s today. The high performance computer market is looking at 100 Gbit links today and will require higher speed links in the future. As most of the mid-range and low-end servers are using 1 Gbit/s links, the migration to 10 Gbit/s links will occur prior to the requirement for 100 Gbit/s.

Some key metrics were repeated themes throughout the meeting. Several server companies discussed the issue of cost of optical links. Within the traditional communications market, most optical component vendors are losing money. The computer industry proposed concentrating on three key metrics: \$/Gbit, \$/W, and comparison to copper. The target of \$1/Gbit is a factor of 4-8 lower than is available today.

Current optical transceiver technology has been developed for the optical communications market. Initially this was a high priced, low volume market. Today it remains at low volume, but with requirements of low price and high reliability. A paradigm shift in optical links will need to occur if the \$1/Gbit or below is to be achieved.

Government initiatives in Japan, Europe, and the United States are pushing for high performance and new technology for optical computing interconnects. Within the U.S., DARPA remains committed to funding new optical technology. The current emphasis on silicon photonics is a key area where potential changes can be realized. The main issue will remain the development of a low cost reliable light source. III-V photonics offers this today, but it is felt incompatible with current silicon industry direction. Fitting III-V into a CMOS environment provides several challenges and a commitment from silicon companies. A solution discussed at the meeting was the potential development of a silicon photonics foundry program, similar to MOSIS. This might enable the silicon photonics paradigm shift to increase research and development in this area.

There several different options today to produce optical circuit boards. The level of sophistication and maturity in polymer-based waveguides has enabled embedded rigid and flex optical circuit cards. The polymer material can be produced in sheets and meet IPC-650 and GR1209 standards. Two critical areas remain to be addressed: coupling out of the PCB and the packaged IC optical interface. MT ferrule connectors offer a solution for card edge connection, but highly tolerant surface coupling remains a packaging concern.

To resolve this, the IC optical interface needs several questions answered. There appears to be no common standard or thinking on how this can occur. Several approaches include low cost VCSELs, while others look to 1310 nm silicon photonic devices to provide the solution.

The IC industry and copper interconnect manufacturers remain committed to pushing the boundary for electrical interconnects. The roadmap for electrical connectors currently extends to 40 Gbit/s. The introduction of silicon via interconnects and 3-D stacking of chipsets will extend the life of on-chip electrical interconnects. The current global wire interconnect problem and “red brick wall” highlighted by the ITRS remain to be resolved. It is evident that the optical technology today for intra-chip and inter-chip signaling remains too immature for introduction. The implementation of silicon photonics into traditional communications space, where they can add increased functionality for smaller package size, does not follow the current trend or philosophy of switch/router manufacturers. Companies, like Cisco Systems, are pulling out CDR technology from the transceivers to reduce cost. Marketing of increased transceiver functionality is a gamble for silicon photonics companies.

While the traditional fiber optics industry is improving and volumes are increasing, the issue remains that the components industry is economically unhealthy. The funding of new interconnect technology is continuing but at a lower level compared to 10 years ago. Several device questions need to be addressed, including the limitation of current low cost devices such as VCSELs. Perhaps new module concepts and packaging are required to open up the server market more aggressively. With continuing price erosion in both the storage and Ethernet markets, the need to develop new technology without a price premium is an industry issue. As Fibre Channel moves from 2xFC to 4xFC to 8xFC, how a company can recoup investment costs is an important question.

Finally, optical interconnects look to have a promising future. As we move forward, the roadmap for this sector needs to be further refined. There remains a need for higher speeds and development of highly reliable, low cost interconnects. The challenge to the industry is twofold: how to fund these requirements and how to standardize the output of the different technologies being implemented today.

Appendix A – Forum Agenda

100Gbit Interconnects and Above: The Need for Speed

February 21, 2007 – Marriott Santa Clara, Santa Clara, CA

7.30 – 8.00 a.m.	Registration and Continental Breakfast
8.00 – 8.10	Welcome – <i>Michael Lebby</i> , OIDA
8.10 – 8.15	Introduction – <i>Bill Ring</i> , WSR Optical Device Solutions
8.15 – 10.00 a.m.	Interconnects: Data Centers
8.15 – 8.35	Hewlett Packard – <i>Terry Morris</i>
8.35 – 8.55	Zarlink – <i>Stan Swirhun</i>
8.55 – 9.15	IBM – <i>Petar Pepeljugoski</i>
9.15 – 9.35	Intel – <i>Tom Willis</i>
9.35 – 9.50	Force10 Networks – <i>John D'Ambrosia</i>
9.50 – 10.00	<i>Moderated discussion</i>
10.00 – 10.20	Coffee Break
10.20 – 12.20 p.m.	Interconnects: Electrical and Optical
10.20 – 10.40	Intel – <i>Jerry Bautista</i>
10.40 – 11.00	TFCG Microsystems Lab, Ghent University – <i>Peter Van Daele</i>
11.00 – 11.20	Force10 Networks – <i>Joel Goergen</i>
11.20 – 11.30	Tyco Electronics – <i>Michael Fogg</i>
11.30 – 11.50	Agilent Technologies – <i>Greg Le Cheminant</i>
11.50 – 12.10 p.m.	Cisco Systems – <i>Mark Gustlin</i>
12.10 – 12.20	<i>Moderated discussion</i>
12.20 – 1.30	Lunch
1.30 – 3.20 p.m.	Interconnects: Active and Passive Optics
1.30 – 1.50	Corning – <i>Richard Grzybowski</i>
1.50 – 2.10	RSoft Design Group – <i>Zhengyu Huang</i>
2.10 – 2.30	Dow Corning – <i>Jon DeGroot</i>
2.30 – 2.50	EMCORE – <i>Kenneth Jackson</i>
2.50 – 3.10	Avago Technologies – <i>Jugnu Ojha</i>
3.10 – 3.20	<i>Moderated discussion; organize breakout sessions</i>

3.20 – 6.00 p.m.	Breakout Discussions
3.20 – 5.00	Breakout and roadmap discussions
5.00 – 5.45	Reports from breakout leaders
5.45 – 6.00	Workshop summary and concluding remarks
6.30 – 8.30 p.m.	Evening Reception

Appendix B – Forum Attendee List

Giovanni Barbarossa

CTO and Sr. Vice President of Product Development
Avanex Corporation
40919 Encyclopedia Circle
Fremont, CA 94538
Tel: (510) 897-4176
Email: Giovanni_Barbarossa@avanex.com

Jerry Bautista

Director, Technology Management
Intel Corporation
SC12-303
2200 Mission College Blvd.
Santa Clara, CA 95054
Tel: (408) 765-8784
Email: jerry.r.bautista@intel.com

John Bowers

Professor
University of California, Santa Barbara
Room 5113, ECE Department
Santa Barbara, CA 93106
Tel: (805) 893-8447
Email: bowers@ece.ucsb.edu

John D'Ambrosia

Scientist, Components Technology
Force10 Networks
7701 Aynlee Way
Harrisburg, PA 17112
Tel: (717) 460-8535
Email: jdambrosia@force10networks.com

Nicola Daldosso

Senior Researcher
University of Trento
via Sommarive 14
Povo (Trento), 38050 Italy
Tel: 390461882037
Email: daldosso@science.unitn.it

Jon DeGroot

Associate Scientist
Dow Corning Corporation
2200 W. Salzburg Road
Midland, MI 48686
Tel: (989) 496-4764
Email: jon.degroot@dowcorning.com

Hongyu Deng

Finisar Corporation
41762 Christy Street
Freemont, CA 94538
Tel: (510) 979-3006
Email: hongyu.deng@finisar.com

Mark Donovan

Senior Director
Finisar Corporation
41762 Christy Street
Sunnyvale, CA 94538
Tel: (510) 979-3002
Email: mark.donovan@finisar.com

Pieter Dumon

Ghent University (Universiteit Gent)
St. Pietersnieuwstraat 41
Dept. of Information Technology
Gent, B-9000
Belgium
Tel: 32-92643448
Email: pieter.dumon@intec.ugent.be

Patrick Ennis

Managing Director
ARCH Venture Partners
1000 Second Avenue, Suite 3700
Seattle, WA 98104
Tel: (206) 674-3273
Email: pjennis@archventure.com

Michael Fogg

Member of the Technical Staff II
Tyco Electronics Corporation
M.S. 128-20, P O. Box 3608
Harrisburg, PA 17105-3608
Tel: (717) 986-5802
Email: mike.fogg@tycoelectronics.com

Veronique Francois

Professor
Ecole de Technologie Superieure
Electrical Engineering Department
1100 Norte-Dame St., West
Montreal, QC, H3C 1K3 Canada
Tel: (514) 396-8800 X7730
Email: vfrancois@ele.etsmtl.ca

Richard Gaughan

Contributing Editor
Photonics Spectra
P. O. Box 1589
Boulder Creek, CA 95006
Tel: (831) 338-1241
Email: rgaughan@mountainoptical.com

Michael Geis

Technical Staff
MIT Lincoln Laboratory
LL-I-300C Lincoln Laboratory
Cambridge, MA 02139
Tel: (781) 981-7875
Email: geis@ll.mit.edu

Andreas Goebel

VP, Photonics Integration
Innolume, Inc.
3333 Bowers Avenue, Suite 190
Santa Clara, CA 95054
Tel: (408) 689-3702
Email: andreas.goebel@innolume.com

Joel Goergen

Chief Scientist/VP of Technology
Force10 Networks
350 Holger Way
San Jose, CA 95134
Tel: (408) 571-3694
Email: joel@force10networks.com

Russ Granger

Business Development Manager
US Conec Ltd.
1555 4th Avenue SE
Hickory, NC 28602
Tel: (828) 267-6385
Email: russgranger@usconec.com

Martin Green

Group Leader, Electronic and Optoelectronic
Materials
NIST
100 Bureau Drive MS 8520
Gaithersburg, MD 20899-8520
Tel: (301) 975-8496
Email: martin.green@nist.gov

Richard Grzybowski

Director, Systems Engineering
Corning Incorporated
One Science Center Drive, SP-AR-02
Corning, NY 14831
Tel: (607) 974-0681
Email: grzybowski@corning.com

Peter Guilfoyle

Chairman, CEO
OptiComp Corporation
215 Elks Point Road, P.O. Box 10779
Zephyr Cove, NV 89448
Tel: (775) 588-4176
Email: peter@opticomp.com

Cary Gunn

CTO
Luxtera
1819 Aston Avenue, Suite 102
Carlsbad, CA 92008
Tel: (760) 448-3520
Email: cary@luxtera.com

Mark Gustlin

Principal Engineer
Cisco Systems, Inc
170 West Tasman Drive
San Jose, CA 95008
Tel: (408) 527-2299
Email: mgustlin@cisco.com

Atikem Haile-Mariam

Product Marketing Manager
Finisar Corporation
1389 Moffett Park Drive
Sunnyvale, CA 94089
Tel: (408) 243-0031
Email: atikem.haile@finisar.com

Sadiq Hasnain

Head, Business Development
National Research Council of Canada
Institute of Microstructural Sciences
1200 Montreal Rd., Building M-50
Ottawa, Ontario, K1A 0R6 Canada
Tel: (613) 993-6040
Email: sadiq.hasnain@nrc-cnrc.gc.ca

Tom Hausken

Director, Optical Components
Strategies Unlimited
201 San Antonio Circle, Suite 205
Mountain View, CA 94040
Tel: (650) 941-3438 Ext. 28
Email: thausken@strategies-u.com

Pierre Herve

Director Enterprise
Intel Corporation
TB1, 8674 Thornton Avenue
Newark, CA 94560
Tel: (510) 578-5889
Email: pierre.m.herve@intel.com

Zhengyu Huang

VP, International Sales and Business
Development
RSoft Design Group
400 Executive Boulevard, Ste. 100
Ossining, NY 10562
Tel: (914) 923-2164
Email: zhengyu@rsoftdesign.com

Bernd Huebner

OSA Development Manager
Finisar Corporation
41762 Christy Street
Sunnyvale, CA 94538
Tel: (408) 203-9308
Email: bernd.huebner@finisar.com

David Huff

VP, Marketing and Business Development
OIDA
1015 N. Trooper Road
Eagleville, PA 19403
Tel: (202) 741-9352
Email: huff@oida.org

Diana Huffaker

Associate Professor
University of New Mexico
Center for High Technology Materials
1313 Goddard, SE
Albuquerque, NM 87106
Tel: (505) 272-7845
Email: huffaker@chtm.unm.edu

David Iams

OIDA
1133 Connecticut Avenue, NW
Suite 600
Washington, DC 20036-4329
Tel: (202) 785-4426
Email: iams@oida.org

Osamu Ibarabi

Senior Engineer
NTT Advanced Technology Corporation
Neocity Mitaka Bldg.
3-35-1 Shimo-renjaku
Mitaka-shi 181-0013, Japan
Tel: 81 29 287 7010
Email: osamu.ibaragi@ntt-at.co.jp

Hideo Itoh

Nanoelectronics Research Institute
National Institute of Advanced Industrial
Science and Technology
AIST Tsukuba Central 2
1-1-1 Umezono
Tsukuba, 305-8568, Japan
Tel: 81298615622
Email: hideo.itoh@aist.go.jp

Kenneth Jackson

EMCORE Fiber Optics
5314 Connemara Drive, NE
Rochester, MN 55906
Tel: (505) 349-0922
Email: kenneth_jackson@emcore.com

John Jaeger

Manager
Infinera
169 Java Drive
Sunnyvale, CA 94089
Tel: (408) 716-4812
Email: jjjaeger@infinera.com

Bahram Jalali

University of California, Los Angeles
6531 Boelter Hall, Box 951595
Material Science & Engineering Dept.
Los Angeles, CA 90095-1595
Tel: (310) 825-4052
Email: jalali@ee.ucla.edu

Siegfried Janz

Group Leader, Optoelectronics Devices
National Research Council of Canada
1200 Montreal Road, Building M-50
Ottawa, ON K1A 0R6 Canada
Tel: (613) 990-0926
Email: siegfried.janz@nrc-cnrc.gc.ca

Pawan Kapur

Research Associate
Stanford University
CIS 100
Stanford, CA 94305
Tel: (650) 906-0844
Email: kapurp@stanford.edu

Jeffrey Kash

Manager and Research Staff Member
IBM Research
T.J. Watson Research Center - Room 29-149
P.O. Box 218
Yorktown Heights, NY 10598
Tel: (914) 945-1448
Email: jeffkash@us.ibm.com

Kenneth Kaufmann

Marketing
Hamamatsu Corporation
360 Foothill Road, Box 6910
Bridgewater, NJ 08807-0910
Tel: (908) 231-0960
Email: kkaufmann@hamamatsu.com

Chris Keller

Manager
Ibiden USA R&D
970 Knox Street, Suite A
Torrance, CA 90502
Tel: (310) 768-0519
Email: ckeller@ibiden-usa.com

Peter Kirkpatrick

Chief Architect
Aprius Inc.
440 North Wolfe Road
Sunnyvale, CA 94085
Tel: (408) 524-3166
Email: peter.kirkpatrick@aprius.com

Ashok Krishnamoorthy

Distinguished Engineer
CTO Physical Science Center
RAS Computer Analysis Laboratory
Sun Microsystems
Mailstop USAN10-107
9515 Towne Centre Drive
San Diego, CA 92121
Tel: (858) 526-9414 Ext. 55414
Email: ashok.k@sun.com

Victor Krutul

Director, Silicon Photonics Strategy
Intel Corporation
2200 Mission College Blvd., SC 12-326
Santa Clara, CA 95052
Tel: (408) 353-3333
Email: victor.krutul@intel.com

Greg Le Cheminant

Measurement Applications Engineer
Agilent Technologies
5301 Stevens Creek Blvd.
Santa Clara, CA 95051
Tel: (707) 577-6524
Email: greg_lecheminant@agilent.com

Michael Lebby

President and CEO
OIDA
1133 Connecticut Avenue, NW
Suite 600
Washington, DC 20036-4329
Tel: (202) 785-4426
Email: lebby@oida.org

Ansheng Liu

Principal Engineer
Intel Corporation
2200 Mission College Blvd.
Santa Clara, CA 95054
Tel: (408) 765-0271
Email: ansheng.liu@intel.com

Jean-Louis Malinge

President and CEO
Kotura
2630 Corporate Place
Monterey Park, CA 91754
Tel: (626) 236-4500
Email: jlmalinge@kotura.com

Jan Meise

Director of Strategic Marketing, Optics
Division
Finisar Corporation
1389 Moffett Park Drive
Sunnyvale, CA 94089
Tel: (408) 548-1000
Email: jan.meise@finisar.com

Terry Morris

Fellow
Hewlett-Packard
3000 Waterview Parkway
Richardson, TX 75080
Tel: (972) 497-4517
Email: terrel.morris@hp.com

Gary Moskovitz

President and Chief Executive Officer
Reflex Photonics Inc.
550 Sherbrooke Street West
Suite 680 West Tower
Montreal, QC H3A 1B9 Canada
Tel: (514) 842-5179 Ext. 226
Email: gmoskovitz@reflexphotonics.com

Ed Murphy

Group Chief Technology Officer
JDSU
45 Griffin Road South
Bloomfield, CT 06002
Tel: (860) 243-6642
Email: ed.murphy@jdsu.com

Maziar Nezhad

Graduate Researcher
University of California, San Diego
Tel: (858) 204-3918
Email: maziar@ucsd.edu

KC Noddings

Photonics Eng. 5
Boeing
901 N. Selby Street
W-S25-C572
El Segundo, CA 90245
Tel: (310) 662-8714
Email: kenneth.c.noddings@boeing.com

Jugnu Ojha

Strategic Marketing
Avago Technologies
350 W. Trimble Road, MS 90th
San Jose, CA 95131-1800
Tel: (408) 435-4270
Email: jugnu.ojha@avagotech.com

Prakash

Research Fellow
University of California, Los Angeles
420 Westwood Plaza
Los Angeles, CA 90095
Tel: (310) 206-4554
Email: prakash@ee.ucla.edu

Peter Pepeljugoski

Research Staff Member Optical Comm.
IBM
1101 Kitchawan Road
Yorktown Heights, NY 10598
Tel: (914) 945-3761
Email: petarp@us.ibm.com

Kannan Raj

Product Line Manager
Zarlink Semiconductor
3650 E. Wier Avenue
Phoenix, AZ 85040
Tel: (602) 414-1707
Email: kannan.raj@zarlink.com

Jaime Relej

Vice President Worldwide Sales
ColorChip Inc.
1028 Hollywood Avenue
Oakland, CA 94602
Tel: (510) 336-2447
Email: jaimer@color-chip.com

William Ring

Managing Director
WSR - ODS
2300 Rachel Terrace, Apt. 7
Pinebrook, NJ 07058
Tel: (908) 212-1923
Email: wsring@wsr-ods.com

Kent Rochford

Optoelectronics Division Chief
NIST
325 Broadway, MS 815
Boulder, CO 80305-3328
Tel: (303) 497-5375
Email: rochford@boulder.nist.gov

Haisheng Rong

Senior Research Scientist
Photonics Technology Lab
Intel Corporation
Corporate Technology Group
2200 Mission College Blvd., SC12-326
Santa Clara, CA 95054
Tel: (408) 765-0331
Email: haisheng.rong@intel.com

Vijit Sabnis

Senior Staff Scientist
Translucent, Inc
952 Commercial Street
Palo Alto, CA 94303
Tel: (650) 213-9311 x225
Email: vijit@translucentinc.com

Scott Schube

Strategic Marketing Manager
Intel Corporation
Optical Platform Division
8674 Thornton Avenue
Newark, CA 94560
Tel: (510) 578-5815
Email: scott.a.schube@intel.com

Mike Shahine

Principal Engineer
Ciena
920 Elkrige Landing Road
Linthicum, MD 21090-2917
Tel: (410) 865-8602
Email: mshahine@ciena.com

Masayuki Shigematsu

President
Innovation Core SEI, Inc.
2580 North First Street, Suite 400
San Jose, CA 95014
Tel: (408) 324-1804
Email: shigematsu-masayuki@sei.co.jp

Etsuji Sugita

Vice President
Hakusan Manufacturing Company
26-5, 2-chome
Minamiikebukuro, Toshima-ku
171-0022 Tokyo, Japan
Tel: 81 3 595 11211
Email: sugita@hakusan-mfg.co.jp

Michael Tan

Scientist
Hewlett-Packard
1501 Page Mill Road, MS 1123
Palo Alto, CA 94304
Tel: (650) 857-5794
Email: mike.tan@hp.com

Marek Tlalka

Vice President of Marketing
Luxtera
1819 Aston Avenue, Suite 102
Carlsbad, CA 92008
Tel: (760) 448-3520
Email: mtlalka@luxtera.com

Hugues Tournier

HW Advisor
Nortel Networks
3500 Carling Avenue
Ottawa, ON K2C 8H9 Canada
Tel: (613) 765-4748
Email: tournier@nortel.com

Peter Van Daele

Professor, Optical Interconnections &
Packaging Group
Ghent University (Universiteit Gent)
INTEC Department
Gaston Crommenlaan 8 Box 201
Gent - Ledenberg
Belgium, B-9050
Tel: 32-9-331-49-04
Email: peter.vandaele@intec.ugent.be

Yuri Vandyshev

Senior Optical Engineer
Finisar Corporation
41762 Christy Street
Fremont, CA 94538
Tel: (510) 979-3044
Email: yuri.vandyshev@finisar.com

Rich Vodhanel

Research Manager, Optical Systems and Networks
Corning Incorporated
1 River Front Plaza
SP-AR-02-4
Corning, NY 14831
Tel: (607) 974-2401
Email: vodhanelrs@corning.com

Guangxi Wang

Graduate Student
CALTECH
Caltech MS 136-93
Pasadena, CA 91125
Tel: (626) 395-2269
Email: wangg@caltech.edu

Winston Way

CTO and Founder
OpVista
870 N. McCarthy Blvd.
Milpitas, CA 95035
Tel: (949) 923-1117
Email: wway@opvista.com

Jan Weem

Principal Engineer
Aprius Inc.
440 North Wolfe
Sunnyvale, CA 94085
Tel: (408) 507-5217
Email: jan.peeters.weem@aprius.com

David Welch

Chief Technology and Strategy Officer
Infinera
169 W. Java Drive
Sunnyvale, CA 94089
Tel: (408) 572-5364
Email: dwelch@infinera.com

Jim Weldon

Chief Operating Officer
Translucent, Inc
952 Commercial Street
Palo Alto, CA 94303
Tel: (650) 213-9311
Email: jim@translucentinc.com

Charles White

Vice President, Strategic Relations
Tessera, Inc.
3099 Orchard Drive
San Jose, CA 95134
Tel: (408) 952-4343
Email: cwhite@tessera.com

Michael Wiemer

PhD Student
Stanford University
P. O. Box 17416
Stanford, CA 94309
Tel: (650) 725-2774

Tom Willis

GM, Optical Interconnects
Intel Corporation
Tel: (503) 957-7567
Email: tom.willis@intel.com

Gregory Wojcik

VP, Engineering
Innolume, Inc.
3333 Bowers Ave., Ste. 190
Santa Clara, CA 95054
Tel: (408) 689-3710
Email: greg.wojcik@innolume.com

Dan-Xia Xu

Senior Research Officer, Optoelectronics Div.
National Research Council of Canada
1200 Montreal Road, Building M-50
Ottawa, ON K1A 0R6 Canada
Tel: (613) 998-8826
Email: Danxia.Xu@nrc-cnrc.gc.ca

Andrew Yang

Consultant
OIDA
3041 N. Pollard St.
Arlington, VA 22207-4137
Tel: (703) 243-2231
Email: acyang@aol.com

Jong-Souk Yeo

Research Scientist
Hewlett-Packard
1000 NE Circle Blvd.
Corvallis, OR 97330
Tel: (541) 715-1641
Email: jong-souk.yeo@hp.com

Homan Yuen

Staff Scientist
Translucent, Inc
952 Commercial Street
Palo Alto, CA 94303
Tel: (650) 213-9311 Ext. 224
Email: homan@translucentinc.com

Eric Zbinden

Director
Aprius Inc.
440 North Wolfe Road
Sunnyvale, CA 94085
Tel: (650) 796-1608
Email: eric.zbinden@aprius.com